

Disambiguating the Role of Noise Correlations When Decoding Neural Populations Together

Hugo Gabriel Eyherabide

*Department of Mathematics and Statistics, Department of Computer Science, and
Helsinki Institute for Information Technology, University of Helsinki
Gustaf Hållströmin katu 2b, FI00560, Helsinki, Finland*

Email: neuralinfo@eyherabidehg.com

Weppage: eyherabidehg.com

Keywords: information theory; neural correlations; independent information; observer perspective; mismatched decoding; brain computations

Abstract

One of the most controversial problems in neural decoding is quantifying the information loss caused by ignoring noise correlations during optimal brain computations. For more than a decade, the measure here called ΔI^{DL} has been believed exact. However, we have recently shown that it can exceed the information loss ΔI^B caused by optimal decoders constructed ignoring noise correlations. Unfortunately, the different information notions underlying ΔI^{DL} and ΔI^B , and the putative rigorous information-theoretical derivation of ΔI^{DL} , both render unclear whether those findings indicate either flaws in ΔI^{DL} or major departures from traditional relations between information and decoding. Here we resolve this paradox and prove that, under certain conditions, observing $\Delta I^{DL} > \Delta I^B$ implies that ΔI^{DL} is flawed. Motivated by this analysis, we test both measures using neural populations that transmit independent information. Our results show that ΔI^{DL} may deem noise correlations more important when decoding the populations together than when decoding them in parallel, whereas the opposite may occur for ΔI^B . We trace these phenomena back, for ΔI^B , to the choice of tie-breaking rules, and for ΔI^{DL} , to unforeseen limitations within its information-theoretical foundations. Our study contributes with better estimates that potentially improve theoretical and experimental inferences currently drawn from ΔI^{DL} without noticing that it may constitute an upper bound. On the practical side, our results promote the design of optimal decoding algorithms and neuroprosthetics without recording noise correlations, thereby saving experimental and computational resources.

1 Introduction

Noise correlations modulate the coactivation of neurons and neural populations at multiple levels in the brain, potentially introducing information that cannot be decoded without knowing their strength and structure. To test this hypothesis, previous studies have assessed, for example, whether correlations are strong, vary across stimuli or experimental tasks, increase the amount of encoded information, or shape spike-triggered averages and decoding filters (Abbott and Dayan, 1999; Brenner et al., 2000; Eyherabide et al., 2008; Gawne and Richmond, 1993; Meister et al., 1995; Nirenberg and Latham, 1998; Panzeri et al., 2001; Quiroga and Panzeri, 2009; Schneidman et al., 2003; Warland et al., 1997). These findings were thought to imply that ignoring noise correlations during decoding must cause an information loss.

This conclusion was challenged by Nirenberg et al. (2001a) who proposed to test the hypothesis directly from the decoder, observer or organism perspective (Bialek et al., 1991; Jaynes, 2003). To that end, they derived an information-theoretical measure of correlation importance here called ΔI^D (also called ΔI , I_{cordep} , ΔI_1 , and ΔI_{NI}^D Eyherabide and Samengo, 2013; Ince et al., 2010; Montani et al., 2007; Nirenberg and Latham, 2003). However, the measure has been perceived as an upper bound, partially because it may potentially exceed the information loss ΔI^B caused by optimal decoders constructed assuming noise independence and even the transmitted information (Eyherabide and Samengo, 2013; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2009, 2010; Schneidman et al., 2003).

These putative limitations were seemingly solved by Latham and Nirenberg (2005) who derived a new information-theoretical measure of correlation importance here called ΔI^{DL} (also denoted ΔI^* and ΔI_{NI}^{DL} in neural decoding, and analogous to Φ^* within integrated information theory of consciousness; Eyherabide and Samengo, 2013; Latham and Nirenberg, 2005; Oizumi et al., 2016). Initially, this measure was employed to justify ΔI^D as an upper-bound of correlation importance (Latham and Nirenberg, 2005). However, its applications have been extended to included the study of medium-to-large populations, thereby overcoming the putative overestimation produced by ΔI^D ; higher-order neural correlations, by combining ΔI^{DL} with maximum-entropy methods; integrated-information theory of consciousness, introducing the decoder perspective and the first measure that seemingly fulfills the theoretical requirements within the field; and neural stochastic codes from the decoding perspective (Eyherabide, 2016; Ince et al., 2010; Latham and Roudi, 2013; Oizumi et al., 2016, 2009, 2010).

Since its introduction, ΔI^{DL} has been regarded as the exact information loss caused by ignoring noise correlations in optimal decoding (Ince et al., 2010;

Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2016, 2009, 2010). However, we have recently shown that, like its predecessor, ΔI^{DL} can also exceed ΔI^B (Eyherabide and Samengo, 2013). Unfortunately, due to the rigorous information-theoretical derivation of ΔI^{DL} and the different information notions underlying ΔI^{DL} and ΔI^B , whether the aforementioned numerical comparison either reveals major departures from traditional views on the relation between information and decoding, or constitutes an indication that ΔI^{DL} is flawed, remains an open question.

To answer this question, we first disentangle the information notion underlying ΔI^{DL} , here called communication information, from the one that we argue underlies ΔI^B , which we call axiomatic information. Taking their differences into account, we determine under which conditions observing that $\Delta I^B < \Delta I^{DL}$ implies that ΔI^{DL} is flawed and overestimates the communication information loss. We also address whether this conclusion can be reached even if ΔI^{DL} does not exceed ΔI^B . To that end, we study neural populations that transmit independent information, and show that ΔI^{DL} grows when decoding them together, as opposed to decoding them in parallel. This paradoxical growth, which can reach about 100 % of transmitted information, is here shown to stem from unforeseen information-theoretical limitations in the derivation of ΔI^{DL} . Surprisingly, we find exactly the opposite phenomena when using ΔI^B , and trace it back to the choice of tie-breaking rules employed during the decoder construction. Our study shows, for the first time, that none of these measures need be additive when information is independent, and most importantly, that ΔI^{DL} need not be exact and may overestimate the communication information loss. Above all, we contribute with tight estimates of communication information losses, thereby potentially improving the accuracy of previous and future theoretical and experimental inferences drawn from them. On the practical side, our results open up new possibilities for simplifying, with tolerable information losses, the computational models that underlie the design of brain-machine interfaces and neuroprosthetics, and for reducing the amount of resources required to study brain computations and information integration (Aflalo et al., 2015; Bouton et al., 2016; Eyherabide and Samengo, 2013; Latham and Nirenberg, 2005; Nirenberg et al., 2001a; Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009; Zhang et al., 2016).

2 Materials and methods

2.1 Notation

Sensory stimuli are here characterized by vectors $\mathbf{S}=[S_1, \dots, S_J]$ of J components, where each j^{th} component S_j represents the value adopted by a different feature. For example, most of our hypothetical experiments employ the following four stimuli: $\boxed{\mathbf{A}}$,

$\textcircled{\text{B}}$, $\textcircled{\text{A}}$ and $\textcircled{\text{B}}$. These stimuli can be characterized using vectors $\mathbf{S}=[S_1, S_2]$ of two components (i.e. $J=2$), where the first component denotes the type of frame (i.e., $S_1 \in \{\square, \bigcirc\}$), and the second component, the type of letter (i.e., $S_2 \in \{A, B\}$).

Neural responses are here characterized by vectors $\mathbf{R}=[\mathbf{R}_1, \dots, \mathbf{R}_J]$ of J components. Each j^{th} component \mathbf{R}_j typically characterizes those aspects of the neural responses (e.g., first-spike latency and spike counts either in the individual or concurrent activity of all neurons; phase and amplitudes in local field potentials or sensor signals from brain-imaging devices; etc.) that are sensitive to the value adopted by the j^{th} stimulus feature. To that end, each \mathbf{R}_j is here characterized as a vector of K_j components, namely $\mathbf{R}_j=[R_j^1, \dots, R_j^{K_j}]$, where each k^{th} component R_j^k denotes the value adopted by the k^{th} response aspect in the set of response aspects that are sensitive to the j^{th} stimulus feature S_j . However, to improve readability, the subscript j will often be eliminated when $J=1$.

As an example, imagine an experiment with two populations of two neurons each, that fire in response to the stimuli mentioned above. The first population is only sensitive to frames and the second population is only sensitive to letters. In this experiment, we will denote the concurrent responses of all neurons as $\mathbf{R}=[\mathbf{R}_1, \mathbf{R}_2]$. Here, \mathbf{R}_1 and \mathbf{R}_2 denote the concurrent responses of all neurons in the first and the second population, respectively. In addition, $\mathbf{R}_1=[R_1^1, R_1^2]$, with R_1^1 and R_1^2 denoting the responses of the first and the second neuron in the first population; \mathbf{R}_2 is defined analogously to \mathbf{R}_1 . Should the four neurons have been sensitive to all stimulus features, we would have denoted their concurrent responses as $\mathbf{R}=[R^1, R^2, R^3, R^4]$.

2.2 Neural encoding

Transforming \mathbf{S} into \mathbf{R} is called encoding (Panzeri et al., 2010). Because the same \mathbf{S} may elicit different \mathbf{R} s, and the same \mathbf{R} may occur for different \mathbf{S} s, both \mathbf{S} and \mathbf{R} are often treated as random variables with joint probabilities $P(\mathbf{S}, \mathbf{R})$. The Shannon or mutual information $I(\mathbf{S}; \mathbf{R})$ encoded in \mathbf{R} about \mathbf{S} is given by the following

$$I(\mathbf{S}; \mathbf{R}) = \sum_{\mathbf{S}, \mathbf{R}} P(\mathbf{S}, \mathbf{R}) \ln \frac{P(\mathbf{S}|\mathbf{R})}{P(\mathbf{S}) P(\mathbf{R})}, \quad (1)$$

where \ln denotes natural logarithms. Consequently, absolute information values are measured in units of nats, as opposed to units of bits (the conversion from nats to bits only requires to divide by $\ln 2$; Cover and Thomas, 2006). However, our choice of natural logarithms does not affect information ratios, and simplifies calculations and notation.

2.3 Noise correlations

The responses of K neurons are deemed noise independent (NI) when the following condition always holds

$$P(\mathbf{R}|\mathbf{S}) = \prod_{k=1}^K P(R^k|\mathbf{S}) \stackrel{\text{def}}{=} P^{NI}(\mathbf{R}|\mathbf{S}). \quad (2)$$

Otherwise, the responses are deemed noise correlated. Here, $\stackrel{\text{def}}{=}$ denotes a definition; and \prod , a product. This definition can be traced back to [Schneidman et al. \(2003\)](#) and takes into account noise correlations at all orders ([Latham and Roudi, 2013](#)), as opposed to those based on linear or nonlinear correlations ([Cohen and Kohn, 2011](#); [Pereda et al., 2005](#)). However, contrary to previous studies ([Delis et al., 2013](#); [Latham and Nirenberg, 2005](#); [Meytlis et al., 2012](#); [Nirenberg and Latham, 2003](#)), here we note that this definition need not be equivalent to those that additionally average across all stimuli ([Gawne and Richmond, 1993](#); [Womelsdorf et al., 2012](#)), which potentially confuse noise correlations with activity correlations ([Schneidman et al., 2003](#)) and are prone to cancellation effects ([Nirenberg and Latham, 2003](#)).

2.4 Neural decoding

Transforming \mathbf{R} into estimated stimuli $\tilde{\mathbf{S}}$ (or into perceptions, decisions and actions) is called decoding ([Panzeri et al., 2010](#)). Analogous to \mathbf{S} and \mathbf{R} , both \mathbf{S} and $\tilde{\mathbf{S}}$ are often treated as random variables with joint probabilities $P(\mathbf{S}, \tilde{\mathbf{S}})$, also called confusion matrix ([Delis et al., 2013](#); [Ince et al., 2010](#); [Quiroga and Panzeri, 2009](#); [Rolls and Treves, 2011](#)). In this study, we focus on optimal decoders, also known as Bayesian or maximum-a-posteriori decoders, ideal homunculus, ideal or Bayesian observers, and optimal-unbiased or maximum-likelihood discrimination ([Bialek, 1987](#); [Ernst and Banks, 2002](#); [Geisler, 2011](#); [Knill and Richards, 1996](#); [Oram et al., 1998](#); [Simoncelli, 2009](#); [Zhang et al., 2016](#)). These decoders map each \mathbf{R} into $\tilde{\mathbf{S}}$ as follows

$$\tilde{\mathbf{S}} = \arg \max_{\mathbf{S}} P(\mathbf{S}|\mathbf{R}), \quad (3)$$

where $P(\mathbf{S}|\mathbf{R})$ is computed via Bayes' rule ([Eyherabide and Samengo, 2013](#)).

When neurons are noise correlated, decoding their concurrent responses \mathbf{R} using the exact $P(\mathbf{R}|\mathbf{S})$ can become experimentally and computationally intractable even for the brain. For this reason, previous studies have proposed to construct optimal decoders assuming that neurons are NI, here called optimal NI decoders ([Eyherabide and Samengo, 2013](#)), but also known as weak-coupling or independent models, weak observers, and naive Bayes classifiers ([Duda et al., 2000](#);

Knill and Richards, 1996; Landy et al., 1995; Meytlis et al., 2012; Nirenberg et al., 2001a; Nirenberg and Latham, 2003; Quiroga and Panzeri, 2009). These decoders map each \mathbf{S} into $\tilde{\mathbf{S}}$ as follows

$$\tilde{\mathbf{S}} = \arg \max_{\mathbf{S}} P^{NI}(\mathbf{S}|\mathbf{R}), \quad (4)$$

with $P^{NI}(\mathbf{S}|\mathbf{R})$ computed from $P^{NI}(\mathbf{R}|\mathbf{S})$ via Bayes' rule.

2.5 Decoding perspective

This study assesses the role of noise correlations from the decoding perspective. Within it, the importance of noise correlations is measured as the losses caused by decoding the actual neural responses assuming that neurons are NI. Noise correlations are deemed important if the losses are significant, and inessential if they are not (see Eyherabide and Samengo, 2013, and references therein). The decoding perspective is often confused with other approaches that measure correlation importance, for example, by comparing the information encoded in real responses with the one encoded in surrogate NI responses (responses generated assuming that neurons are NI). Even when using decoders (Delis et al., 2013; Nirenberg and Latham, 1998; Quiroga and Panzeri, 2009), these approaches need be neither conceptually nor quantitatively related to the decoding perspective (Averbeck et al., 2006; Latham and Nirenberg, 2005; Nirenberg and Latham, 2003). Avoiding such confusion is fundamental to correctly interpreting our results and conclusions.

2.6 Measures of correlation importance

Even within the decoding perspective, the choice of correlation-importance measure remains controversial. Here we will study the relation between the following three commonly-used measures

$$\Delta I^B(S, \mathbf{R}) = \sum_{S, \mathbf{R}} P(S, \mathbf{R}) \ln \frac{P(S|\mathbf{R})}{P(S|\tilde{\mathbf{S}})} \quad (5)$$

$$\Delta I^D(S, \mathbf{R}) = \sum_{S, \mathbf{R}} P(S, \mathbf{R}) \ln \frac{P(\mathbf{S}|\mathbf{R})}{P^{NI}(\mathbf{S}|\mathbf{R})} \quad (6)$$

$$\Delta I^{DL}(S, \mathbf{R}) = \min_{\theta} \sum_{S, \mathbf{R}} P(S, \mathbf{R}) \ln \frac{P(\mathbf{S}|\mathbf{R})}{P^{NI}(\mathbf{S}|\mathbf{R}, \theta)} \stackrel{\text{def}}{=} \min_{\theta} \delta I(\mathbf{S}; \mathbf{R}|\theta), \quad (7)$$

where $\tilde{\mathbf{S}}$ is the output of an optimal NI decoder (Section 2.4), and

$$P^{NI}(\mathbf{S}|\mathbf{R}, \theta) \propto P(\mathbf{S}) P^{NI}(\mathbf{R}|\mathbf{S})^\theta, \quad (8)$$

with the convention that $0^0=0$ for automatically overcoming the drawbacks of previous definitions found in (Eyherabide and Samengo, 2013). These measures have been previously related to the information loss caused by ignoring noise correlations in optimal decoding (Ince et al., 2010; Latham and Nirenberg, 2005; Nirenberg and Latham, 2003). Accordingly, the decoded information when ignoring noise correlations in optimal decoding has previously been quantified by subtracting these measures from the encoded information (Oizumi et al., 2009, 2010). Further details about their derivations and interpretations are here postponed until Section 3, where we will reassess them in the context of the present study.

2.7 Independent information

The notion of independent information has previously been given different definitions that need not be interchangeable. In this study, we say that the responses $\mathbf{R}_1, \dots, \mathbf{R}_J$ of J neural populations transmit independent information when they fire independently and selectively to J independent stimulus features S_1, \dots, S_J (Fano, 1961). Using the notation introduced in Section 2.1, our definition implies the following

$$P(\mathbf{S}, \mathbf{R}) = \prod_{j=1}^J P(S_j, \mathbf{R}_j). \quad (9)$$

When (9) holds, each pair $[S_j, \mathbf{R}_j]$ is here said to constitute an independent-information channel or stream.

Under our definition, the information carried by independent-information streams is additive (Fano, 1961), namely

$$I(\mathbf{S}, \mathbf{R}) = \sum_{j=1}^J I(S_j; \mathbf{R}_j). \quad (10)$$

For simplicity, hereinafter we abbreviate the notation for the arguments of any information measure X as follows

$$X_{1,\dots,J} \stackrel{\text{def}}{=} X(S_1, \dots, S_J; \mathbf{R}_1, \dots, \mathbf{R}_J) \quad (11)$$

$$X_j \stackrel{\text{def}}{=} X(S_j, \mathbf{R}_j) \quad (12)$$

$$X_{1+\dots+J} \stackrel{\text{def}}{=} X_1 + \dots + X_J \quad (13)$$

Using these abbreviations, the additivity of the information carried by independent-information streams can be simply put as $I_{1,\dots,J} = I_{1+\dots+J}$.

This property has often been used as the actual definition of independent information (Brenner et al., 2000; Rolls and Treves, 2011; Schneidman et al., 2003, 2011), whereas other definitions only require that populations are asymptotically or conditionally independent (Cover and Thomas, 2006; Gawne and Richmond, 1993; Samengo and Treves, 2000). Our definition is more stringent than those and ensures that independent information remains independent after arbitrary parallel transformations of the form $[\tilde{S}_j, \tilde{\mathbf{R}}_j] = f_j(S_j, \mathbf{R}_j)$, namely

$$P(\tilde{S}_1, \tilde{\mathbf{R}}_1, \dots, \tilde{S}_J, \tilde{\mathbf{R}}_J) = \prod_{j=1}^J P(\tilde{S}_j, \tilde{\mathbf{R}}_j). \quad (14)$$

This property plays a fundamental role in our study, and it also holds for the definitions given in Cover and Thomas (2006) and Eyherabide and Samengo (2010). However, our definition is more general than those for the noise in each j^{th} stream is allowed to depend on the j^{th} feature.

3 Results

3.1 Underlying information notions need not be reliably related

Quantifying the information loss caused by ignoring noise correlations from the decoding perspective remains controversial, oftentimes due to unfulfilled expectations about the relation between different measures of information loss. One of the most important unfulfilled expectations concerns the measures ΔI^{DL} and ΔI^B (Section 2.6). Because ΔI^{DL} continues to be considered the exact information loss in optimal NI decoding (Section 1 and references therein), it may seem natural to expect that ΔI^{DL} constitute a lower bound on the information loss ΔI^B caused by optimal NI decoders. However, as we have recently shown, this relation need not hold (Eyherabide and Samengo, 2013). In this section, we begin our quest to disentangle whether this result indicates that the measures are flawed or that traditional expectations are unjustified, by comparing the information notions underlying ΔI^{DL} and ΔI^B .

The measure ΔI^{DL} was derived by Latham and Nirenberg (2005) using a notion of information with roots in communication theory, here called communication information (Section 2.2). Within the context of this study, this notion and the derivation of ΔI^{DL} can both be intuitively described using the population of two neurons depicted in Fig. 1(a). These two neurons can distinguish between two visual stimuli, namely \square and \bigcirc , by concurrently firing the same number of spikes (1 or 2) after observing \square , and different number of spikes (2 or 3) after observing \bigcirc .

To compute ΔI^{DL} , Latham and Nirenberg (2005) employed sequences $\mathbf{S}^{[N]} = [\mathbf{S}_1^{[N]}, \dots, \mathbf{S}_N^{[N]}]$ of N independent and identically distributed stimuli, and grouped

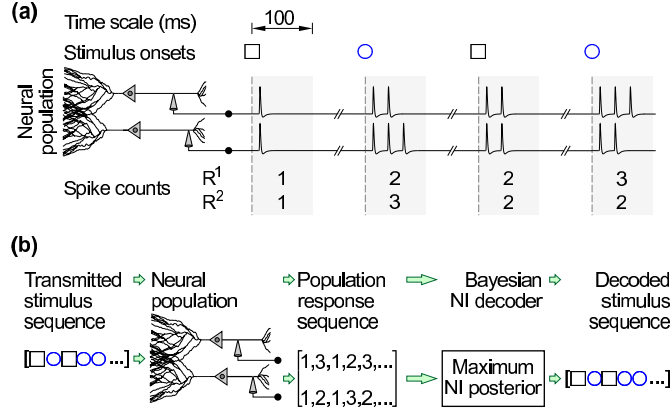


Figure 1: Communication notions previously used to derive ΔI^{DL} . (a) Hypothetical intracellular recording of the concurrent spike trains produced by two neurons within 100 ms after presenting each of the visual stimuli \square and \circ in alternation, with pauses between each presentation. Only two types of population responses occur for each stimulus, and can be fully described using spike counts. (b) Representation of the memoryless communication system employed by Latham and Nirenberg (2005) to compute ΔI^{DL} .

into sets called codebooks (Fig. 1(b)). Each $\mathbf{S}^{[N]}$ is transformed by the aforementioned neural population, one stimulus at a time and independently of the others, into a sequence $\mathbf{R}^{[N]} = [\mathbf{R}_1^{[N]}, \dots, \mathbf{R}_N^{[N]}]$ of N population responses. The $\mathbf{R}^{[N]}$ are then fed into an optimal NI decoder that attempts to reconstruct the original $\mathbf{S}^{[N]}$ (Section 2.4).

Using this interpretation, Latham and Nirenberg (2005) computed an estimate of the average probability $P(\mathbf{S}^{[N]} \neq \tilde{\mathbf{S}}^{[N]})$ that the decoded sequences $\tilde{\mathbf{S}}^{[N]}$ produced by optimal NI decoders differ from the transmitted sequences $\mathbf{S}^{[N]}$, usually denoted $P_e^{[N]}$. They showed this estimate to decay exponentially as N grows for codebooks of up to $\exp[N \tilde{I}]$ sequences, from which they concluded, based on standard information-theoretical results, that \tilde{I} quantifies the communicated information. Subtracting \tilde{I} from the encoded information yielded the measure ΔI^{DL} that is currently believed the exact information loss caused when ignoring noise correlations in optimal decoding (Section 2.6).

Current beliefs notwithstanding, we have recently shown that ΔI^{DL} can exceed the information loss ΔI^B caused by optimal decoders constructed ignoring noise correlations (Section 2.6; Eyherabide and Samengo, 2013). Analogously to previous conclusions about ΔI^D , this finding may seem to directly indicate that ΔI^{DL} is flawed and overestimates the information loss caused by ignoring noise correlations in optimal decoding. However, as we note here, this conclusion overlooks the fact that ΔI^B is fundamentally different from ΔI^{DL} , and therefore not necessarily comparable.

Specifically, ΔI^B differs from ΔI^{DL} in at least the following three fundamental as-

pects. First, it stems from treating decoding as a recoding process, rather than as the final stage in a communication system (Quiroga and Panzeri, 2009). Second, it is sensitive to random errors but insensitive to systematic errors, and hence be large without decoders ever being correct (Quiroga and Panzeri, 2009; Schneidman et al., 2003). Third, its derivation involves single stimuli, as opposed to stimulus sequence and asymptotic limits. These three differences are not necessarily unknown in the neuroscience literature, but they are often overlooked.

Overlooking the differences between ΔI^B and ΔI^{DL} , and therefore deeming them comparable, may seem justified for at least two reasons. First, previous studies have shown that ΔI^D (Latham and Roudi, 2013; Nirenberg and Latham, 2003), and consequently ΔI^{DL} (Eyherabide, 2016), reduces to the traditional information loss when applied to ignoring response aspects through transformations of the population response. Second, the data processing inequality ensures that, under those cases, ΔI^{DL} is a lower bound of the information loss caused by decoders that operates on the transformed responses. However, as we have recently shown, the second reason is invalid because, within the decoding perspective, decoders operate on the original responses, as opposed to the transformed ones. Furthermore, we have recently proved that, contrary to previously thought, the first reason need not be valid even for deterministic transformations (Eyherabide, 2016).

Therefore, we find questionable both to overlook the differences between the information notions underlying ΔI^{DL} and ΔI^B , and to draw conclusions about flaws in ΔI^{DL} based on previous observations that $\Delta I^{DL} > \Delta I^B$. Instead, those observations may indicate major departures from traditional relations between their underlying information notions. Nevertheless, it seems to us rather unfortunate that a putative exact measure of information loss cannot be regarded as a yardstick against which the performance of optimal NI decoders can be measured.

Resolving this puzzle requires that we hereinafter undertake at least two actions. First, distinguishing the notion of communication information underlying ΔI^{DL} , from the notion of axiomatic information underlying ΔI^B , here so-called because it seemingly interprets information as an abstract measure of arbitrary correlations that fulfills certain intuitive and desirable axioms (Fano, 1961; Gallager, 1968; Shannon and Weaver, 1949; Woodward and Davies, 1952). Second, assessing the properties of these two measures within the boundaries of their underlying information notions. In this way, our strategy becomes unfortunately more complex than those followed by previous studies, but guarantees to avoid their potential confounds, and to accurately compare ΔI^B and ΔI^{DL} .

3.2 ΔI^{DL} overestimates the communication information loss

We have recently shown not only that ΔI^{DL} can exceed ΔI^B , but also that ΔI^{DL} can be positive even when optimal NI decoders never make mistakes (Eyherabide and Samengo, 2013). However, the results there shown comprised a single hypothetical experiment analogous to the one depicted in Fig. 1(a), which may seem to possess peculiar characteristics. Most importantly, we did not connect those results with the possibility of flaws in ΔI^{DL} . In this section, we test this possibility, avoiding potential confounds by recasting the above experiment within the framework of Latham and Nirenberg (2005), and generalizing the results to any experiment in which optimal NI decoders never make mistakes, regardless of the stimulus-response distributions.

Imagine generic experiments in which optimal NI decoders can perfectly identify the \mathbf{S} that elicited each \mathbf{R} , including and beyond that in Fig. 1(a). Mathematically, this means that $P(\mathbf{S} \neq \tilde{\mathbf{S}}) = 0$, where $\tilde{\mathbf{S}}$ denotes the decoded stimuli. As we note here, all these experiments can be recast within the framework of Latham and Nirenberg (2005), by interpreting neural populations as memoryless channels that read unit-length stimulus sequences $\mathbf{S}_1^{[N]}$. These sequences are turned into unit-length population-response sequences $\mathbf{R}_1^{[N]}$, which are subsequently fed into optimal NI decoders.

Suppose now that the length of the sequences is increased to an arbitrary value N . According to ΔI^{DL} , optimal NI decoders may start making mistakes. Indeed, suppose that, in the experiment of Fig. 1(a), the population responses $\mathbf{R} = [R^1, R^2]$ associated with each stimulus are equally-likely and \square occurs more frequently than \bigcirc . Mathematically, this implies that $P(\mathbf{R}|\mathbf{S}) = 0.5$ regardless of \mathbf{R} and \mathbf{S} , and that $P(\square) > 0.5$. In that case, ΔI^{DL} is positive and can reach $\approx 25\%$ of the transmitted information (Eyherabide and Samengo, 2013), both of which can only occur if $P_e^{[N]} > 0$.

To prove that this is not the case, recall that by hypothesis, optimal NI decoders can produce stimulus estimates $\tilde{\mathbf{S}}_n^{[N]}$ for each n^{th} population response $\mathbf{R}_n^{[N]}$ within the received $\mathbf{R}^{[N]}$ without errors, namely $P(\mathbf{S}_n^{[N]} \neq \tilde{\mathbf{S}}_n^{[N]}) = 0$. Boole's inequality (Casella and Berger, 2002) ensures that concatenating these parallel estimates to produce the decoded stimulus sequence $\tilde{\mathbf{S}}^{[N]}$ yields no sequence-errors, that is

$$P(\mathbf{S}^{[N]} \neq \tilde{\mathbf{S}}^{[N]}) \leq N \max_n P(\mathbf{S}_n^{[N]} \neq \tilde{\mathbf{S}}_n^{[N]}) = 0. \quad (15)$$

Following Latham and Nirenberg (2005), this result implies that, in the above experiments, ignoring noise correlations in optimal decoding never causes communication information losses.

Therefore, we have proved for the first time that ΔI^{DL} is flawed and overestimates the communication information loss when ignoring noise correlations in optimal decod-

ing. For this reason, we denote the latter as ΔCI . Momentarily, we will define ΔCI as follows

$$\Delta CI = \begin{cases} 0 & \text{in the absence of decoding errors} \\ \Delta I^{DL} & \text{otherwise} \end{cases}, \quad (16)$$

but refine it later after revealing additional sources of overestimation.

3.3 Misleading intuitions and limitations

Although we proved above that ΔI^{DL} overestimates ΔCI , our proof still partially rests on the experiment in Fig. 1(a). This experiment possesses three properties that may seem peculiar and in direct contradiction with current intuitions about the role of noise correlations in optimal decoding. In this section, we show these properties inessential for ΔI^{DL} to overestimate ΔCI .

Our demonstrations focus on cases in which optimal NI decoders make no mistakes. This property implies that each \mathbf{R} occurs for only one \mathbf{S} . However, we exclude the cases in which each \mathbf{R} would occur for only one \mathbf{S} should neurons be NI, because the resulting ΔI^{DL} is trivially zero and therefore tight. In the remaining cases, at least one \mathbf{R} that occurs when neurons are correlated always occur for more than one \mathbf{S} should neurons be NI.

To illustrate this condition, we constructed three experiments analogous to Fig. 1(a) and represented them using Cartesian coordinates (Fig. 2). In all panels, the \mathbf{R} s recorded during the experiment are represented at the top, and those that would have been recorded should neurons be NI are represented at the bottom. The experiment in Fig. 2(a) actually coincides with that in Fig. 1(a), whereas those in Fig. 2(b) and Fig. 2(c) are variations of it explained below. In all cases, the Cartesian representations clearly show that each \mathbf{R} occurs for only one \mathbf{S} when neurons are correlated (top) and that $\mathbf{R}=[2, 2]$ would have occurred for all \mathbf{S} s should neurons be NI.

The first seemingly peculiar property of Fig. 1(a) is that noise correlations vary across \mathbf{S} . Indeed, Fig. 2(a) shows that the concurrent responses of both neurons are positively correlated when elicited by \square (that is, the Pearson correlation coefficient $\rho(R^1, R^2|\square) > 0$), and negatively correlated when elicited by \bigcirc (that is, $\rho(R^1, R^2|\bigcirc) < 0$). This observation has previously been thought to indicate that noise correlation must be important from the decoding perspective (Latham and Nirenberg, 2005; Nirenberg et al., 2001b). However, in Section 3.2, we proved that this intuition need not be correct.

Most importantly, this first property is insufficient and unnecessary for proving that ΔI^{DL} overestimates ΔCI . To prove it insufficient, we take the experiment in

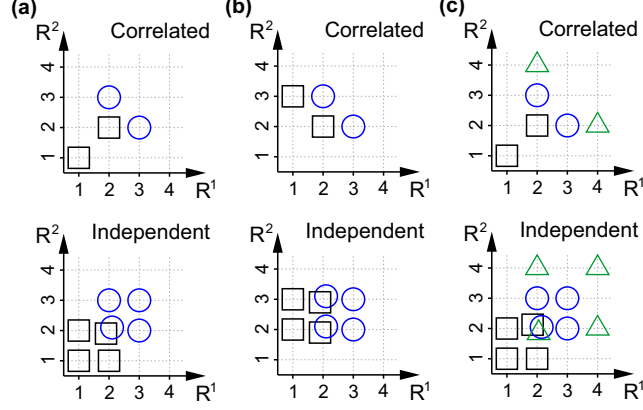


Figure 2: ΔI^{DL} may overestimate the communication information loss. (a) Cartesian representation of the responses recorded in the experiment of Fig. 1(a) (top) and the surrogate NI responses that would have occurred should neurons be NI (bottom). (b) Analogous description to panel (a), but for the responses recorded in another hypothetical experiment. (c) Analogous description to panel (b). In all cases, the symbols represent the stimuli; and their locations, the concurrent responses they elicit. To improve visibility, symbols that occur in the same location are shown to only partially overlap.

Fig. 2(a) and set $P(\mathbf{R}|\mathbf{S}) \neq 0.5$ regardless of \mathbf{S} and \mathbf{R} , thereby obtaining through (7) that $\Delta I^{DL} = 0$ (Eyherabide and Samengo, 2013). To prove it unnecessary, we build another experiment analogous to Fig. 2(a), but with $\mathbf{R} = [1, 1]$ replaced by $[1, 3]$ (Fig. 2(b)). Contrary to Fig. 2(a), here noise correlations remain constant across stimuli both in sign and in strength (i.e., $\rho(R^1, R^2|\square) = \rho(R^1, R^2|\circ) = -1$). Notwithstanding, setting $P(\square) = 0.52$, $P(\mathbf{R} = [2, 2]|\square) = 0.48$, and $P(\mathbf{R} = [2, 3]|\circ) = 0.52$, yields $\Delta I^{DL} \approx 50\%$ even though optimal NI decoders can be readily proved to never make mistakes.

The second seemingly peculiar property is that the value of ΔI^{DL} equals the value of its predecessor, namely ΔI^D (Eyherabide and Samengo, 2013). Interestingly, this observation has previously been thought to indicate that ΔI^D is tight (Latham and Roudi, 2013). On the contrary, here we show that such observation actually indicates that ΔI^{DL} is loose. Specifically, we can prove that $\Delta I^{DL} = \Delta I^D$ is sufficient but inessential for proving that ΔI^{DL} overestimates the communication information loss. To avoid clutter, we hereinafter omit the response labels from the arguments of the probabilities.

To prove it sufficient, notice that, for the cases under consideration (see second paragraph), at least one \mathbf{R} and \mathbf{S} exists for which $P^{NI}(\mathbf{S}|\mathbf{R}) < 1$ and $P(\mathbf{S}|\mathbf{R}) = 1$. Therefore, (6) yields that $\Delta I^D > 0$, and consequently, the second property implies that $\Delta I^{DL} > 0$. However, the second property need not imply that optimal NI decoders make mistakes. Indeed, the property holds when setting $P(\mathbf{S}, \mathbf{R})$ in Fig. 2(a) so that $P^{NI}([2, 2]|\square) = P^{NI}([2, 2]|\circ)$, even though optimal NI decoders never make mistakes

when $P(\square) > 0.5$.

To prove it inessential, we set in Fig. 2(b) $P(\square) = 0.48$, $P([2, 2]|\square) = 0.58$, and $P([2, 3]|\bigcirc) = 0.48$. Using these probabilities, (6) and (7) readily yield different values for ΔI^D ($\approx 48\%$) and ΔI^{DL} ($\approx 36\%$), respectively. However, optimal NI decoders can be readily proved to never make mistakes.

The third seemingly peculiar property is that each \mathbf{R} that would occur for more than one \mathbf{S} should neurons be NI, would do so with equal frequency for those \mathbf{S} . Mathematically, $P^{NI}(\mathbf{R}|\mathbf{S}_A) > 0$ and $P^{NI}(\mathbf{R}|\mathbf{S}_B) > 0$ implies that $P^{NI}(\mathbf{R}|\mathbf{S}_A) = P^{NI}(\mathbf{R}|\mathbf{S}_B)$. However, this property is sufficient but typically inessential for proving that ΔI^{DL} overestimates ΔCI .

We prove it sufficient by noting that this property turns $\delta I(\mathbf{S}; \mathbf{R}|\theta)$ independent of θ , and thus $\Delta I^{DL} = \Delta I^D$. To prove it typically inessential, we first note that the third property is necessary for experiments comprising only two \mathbf{S} and only one \mathbf{R} for which $P^{NI}(\mathbf{R}|\mathbf{S}_A) = P^{NI}(\mathbf{R}|\mathbf{S}_B)$, because otherwise $\delta I(\mathbf{S}; \mathbf{R}|\theta) \rightarrow 0$ either when $\theta \rightarrow -\infty$ or $+\infty$, thereby yielding $\Delta I^{DL} = 0$.

However, this need not be the case when experiments comprise either more than one \mathbf{R} for which $P^{NI}(\mathbf{R}|\mathbf{S}_A) = P^{NI}(\mathbf{R}|\mathbf{S}_B)$ even if for only two stimuli, as we have already shown during our proofs of the first and the second properties, or a single \mathbf{R} with $P^{NI}(\mathbf{R}|\mathbf{S}_A) = P^{NI}(\mathbf{R}|\mathbf{S}_B)$ for more than two stimuli. To prove the latter, we take the experiment of Fig. 2(a) and add two population responses elicited by a third stimulus \triangle (Fig. 2(c)). Setting $P(\mathbf{R}|\mathbf{S})$ so that $P^{NI}([2, 2]|\triangle) < P^{NI}([2, 2]|\square) < P^{NI}([2, 2]|\bigcirc)$ (or in reversed order) turns $\Delta I^{DL} > 0$ regardless of the stimulus probabilities, even though it can be readily proved that stimulus probabilities always exist for which optimal NI decoders never make mistakes.

To summarize, we have shown that the seemingly peculiar properties of the experiment in Fig. 1(a) are inessential for proving that ΔI^{DL} overestimates ΔCI . Contrary to previous beliefs, we have also found that the sign and strength of noise correlations can be misleading about their role in optimal decoding. Most importantly, we have proved that the proximity of ΔI^D to ΔI^{DL} need not indicate that ΔI^D is close to ΔCI , but that ΔI^{DL} is loose.

3.4 Overestimation in the presence of decoding errors

Our analysis so far may seem limited to cases in which optimal NI decoder make no mistakes, which is arguably rare in the nervous system. In this section, we show that the measure ΔI^{DL} may overestimate ΔCI even when optimal NI decoders do make mistakes. To prove this, consider the experiment shown in Fig. 3(a) depicting the concurrent responses $\mathbf{R} = [R_1^1, R_1^2, R_2^1]$ of three neurons elicited by four different \mathbf{S} , namely

$\textcircled{\text{A}}$, $\textcircled{\text{B}}$, $\boxed{\text{A}}$ and $\boxed{\text{B}}$.

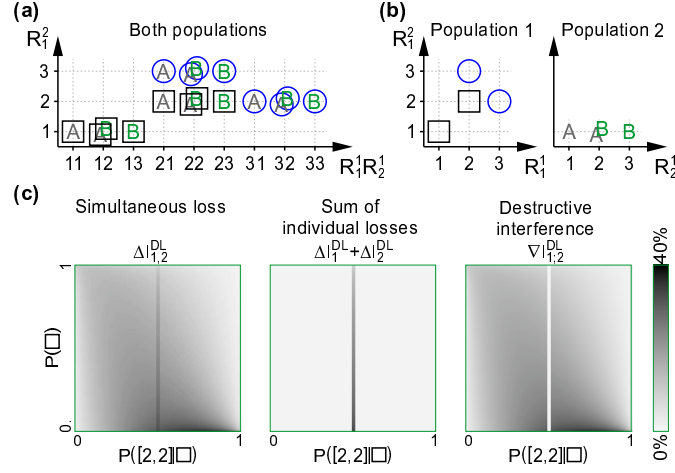


Figure 3: Correlations become paradoxically important when integrating independent information. (a) Cartesian representation of the concurrent responses $\mathbf{R}=[R_1^1, R_1^2, R_2^1]$ of three neurons elicited by four different compound stimuli ($\textcircled{\text{A}}$, $\textcircled{\text{B}}$, $\boxed{\text{A}}$ and $\boxed{\text{B}}$). (b) Cartesian representations of the concurrent responses of R_1^1 and R_1^2 (left), which are sensitive only to variations in the frames; and the responses R_2^1 (right), which are only sensitive to variations in the letters. (c) Density or heat maps (with color-scale given on the right) depicting, as a function of $P(\square)$ and $P([2,2]|\square)$, the following three quantities: the value of ΔI^{DL} when decoding frames and letters together (left); the sum of the values of ΔI^{DL} when decoding them singly (middle); and the resulting destructive interference (right). The line $P([2,2]|\square)=0.5$ has been widened to ease the visualization of discontinuities.

In this experiment, $\mathbf{S}=[S_1, S_2]$ consist of two independently chosen features: a frame ($S_1 \in \{\square, \bigcirc\}$) and a letter ($S_2 \in \{\text{A}, \text{B}\}$). The neurons can be separated into two populations that fire independently of one another and selectively to different stimulus features (Fig. 3(b)). The first population is only sensitive to the frames and analogous to that in Fig. 1(a). The second population is only sensitive to the letters, and consists of a single neuron. As a result, each population constitutes an independent information stream, namely $P(\mathbf{S}, \mathbf{R})=P(S_1, \mathbf{R}_1) P(S_2, \mathbf{R}_2)$ (Section 2.7).

Suppose that we set for the second population $P(2|\text{A})=P(2|\text{B})$. In that case, (7) readily yields that the value of ΔI^{DL} computed using all neurons equals the one computed using only the first two. However, these two neurons are analogous to those in Fig. 1(a), for which we have already shown that ΔI^{DL} may overestimate ΔCI . Therefore, it seems evident that ΔI^{DL} also overestimates ΔCI for the experiment in Fig. 3(a) as well.

Using this result, we can refine the computation of ΔCI as follows

$$\Delta CI \begin{cases} = 0 & \text{if } P(\mathbf{S} \neq \tilde{\mathbf{S}}) = 0 \\ \leq \Delta I^{DL} & \text{otherwise} \end{cases}, \quad (17)$$

thereby reflecting our findings that, contrary to previous studies, ΔI^{DL} cannot be ensured tight. Although this conclusion is correct, the rationale need not be general and should therefore be observed with caution. The reason for our concerns lies in paradoxical properties of ΔI^{DL} that, as we reveal in the next section, emerge when setting $P(\mathbf{S}, \mathbf{R})$ to other values than those used in our demonstration.

3.5 The whole is less than the sum of its independent parts

We have just shown for the first time that ΔI^{DL} may overestimate ΔCI , but the overestimation may seem to occur only when at least some stimulus features (as opposed to stimulus identities) can be perfectly identified. These cases are arguably rare in the nervous system where stimulus features can almost never be perfectly identified due to noise, in which case ΔI^{DL} may still be exact. Surprisingly, here we show that, in those cases, the value of ΔI^{DL} computed over all neural populations typically exceeds the sum of the values of ΔI^{DL} computed over each neural population, even when neural populations fire independently and selectively to different and independent stimulus features, thereby transmitting independent information (section 2.7).

Specifically, suppose that in Fig. 3(a) we set $P(\mathbf{B})$ equal to $P(\square)$. In addition, we set $P(1|\mathbf{A})$, $P(2|\mathbf{B})$, and $P([2, 3]|\bigcirc)$, equal to $P([2, 2]|\square)$, where without risk of confusion we have omitted the response labels R_j^k within the arguments of the probabilities to avoid clutter. Under these conditions, the computation of ΔI^{DL} using all neurons in both populations yields a value here denoted $\Delta I_{1,2}^{DL}$ given by the following

$$\Delta I_{1,2}^{DL} = \begin{cases} -0.5 q \ln q & \text{if } \alpha = 0.5 \\ \phi \mathcal{H}(q\alpha/\phi) - q\alpha \ln 4 & \text{otherwise,} \end{cases} \quad (18)$$

where for compactness, we have defined $q = P(\square)$, $\alpha = P([2, 2]|\square)$, $\mathcal{H}(x) = -x \ln x - \bar{x} \ln \bar{x}$ (also called binary entropy function), $\phi = 2q\alpha + \bar{q}\bar{\alpha}$, and the bar over a symbol denotes complement to unity (e.g., $\bar{z} = 1 - z$).

Notice that Pearson correlation coefficients (ρ) cannot fully characterize the response distributions. Indeed, the only nonzero ρ s arise for the responses of both neurons in population 1. For \mathbf{A} and \mathbf{B} , $\rho = 1$, whereas for \mathbf{A} and \mathbf{B} , $\rho = -1$, and their values remain constant regardless of $P([2, 2]|\square)$. In addition, ΔI^{DL} depends on $P([2, 2]|\square)$, as we show in (18) and Fig. 3(c), thereby rendering ρ also unsuitable for assessing the properties of ΔI^{DL} .

The value of $\Delta I_{1,2}^{DL}$ in Fig. 3(a) is almost always positive, and can reach $\approx 100\%$ of the transmitted information when $P(\square) \rightarrow 0$ and $P([2, 2]|\square) \rightarrow 1$ (Fig. 3(c) left, and ??). According to previous studies, this result would indicate that noise correlations are important in optimal decoding or even crucial (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2016, 2010). However, we find the above result paradoxical because noise correlations only exist within population 1, and, unless $P([2, 2]|\square) = 0.5$, they are irrelevant for decoding that population (Eyherabide and Samengo, 2013).

Indeed, the computation of ΔI^{DL} using only the neurons in the first population always yields $\Delta I_1^{DL} = 0$, unless $P([2, 2]|\square) = 0.5$, whereas its computation using only the neurons in the second population always yields $\Delta I_2^{DL} = 0$, regardless of $P([2, 2]|\square)$ (Fig. 3(c) middle). Therefore, the positive values of $\Delta I_{1,2}^{DL}$ seem to arise from a paradoxical growth in correlation importance caused by decoding the populations together, as opposed to decoding them in parallel, even though the populations transmit independent information. This paradoxical growth is here called destructive interference, and quantified as the following difference

$$\nabla I_{1,2}^{DL} = \Delta I_{1,2}^{DL} - \Delta I_1^{DL} - \Delta I_2^{DL}. \quad (19)$$

In our experiment, $\nabla I_{1,2}^{DL}$ is always positive and equal to $\Delta I_{1,2}^{DL}$ (though zero when $P([2, 2]|\square) = 0.5$), and therefore can reach $\approx 100\%$ of the transmitted information (Fig. 3(c) right).

Notice that the above results have not been corrected for the overestimation problems described in Section 3.2 and Section 3.4. As a result, the observed destructive interference cannot be directly attributed to an overestimation of the types there studied. Instead, it constitutes a new phenomena that can occur in parallel with other sources of overestimation.

Indeed, as shown in Section 3.2, the actual communication information loss ΔCI_1 computed using only the neurons in the first population is equal to zero when optimal NI decoders make no mistakes, and this can occur even if $P([2, 2]|\square) = 0.5$. To that end, one of the following two conditions must hold: $P(\square) > 0.5$, or $P(\square) = 0.5$ but optimal NI decoders are constructed using tie-breaking rules that choose \square whenever $P^{NI}(\square|\mathbf{R}_1) = P^{NI}(\bigcirc|\mathbf{R}_1)$ (Eyherabide and Samengo, 2013). Unfortunately, we can not refine the estimation of the actual communication information loss $\Delta CI_{1,2}$ computed using all neurons in all populations based on the results in Section 3.4 because, unlike the example there studied, here $\Delta I_{1,2}^{DL} > \Delta I_1^{DL}$. This result would imply that the actual destructive interference is zero only when $P([2, 2]|\square) = 0.5$ and $P(\square) < 0.5$, that is, on the lower half of the line defined by $P([2, 2]|\square) = 0.5$ in Fig. 3(c).

In conclusion, here we have shown for the first time that ΔI^{DL} is superadditive

when information is independent. This result is in stark contrast with the additivity of its predecessor, ΔI^D , which we later prove in [Section 3.7](#). Most importantly, our observation reveals a major departure from traditional views on the desirable properties of information measures ([Cover and Thomas, 2006](#); [Fano, 1961](#)). From a neuroscientific perspective, our result implies that, should ΔI^{DL} be exact as currently thought, noise correlations would paradoxically become more important when analyzing neural populations as a whole than when analyzing each of its constituent parts, even when the parts transmit independent information.

3.6 Destructive interference in brain models

Information in the brain may be transmitted by one or multiple heterogeneous neural substrates (e.g. single neurons, neural populations, cortical areas or functional networks) using different temporal scales, frequency bands, amplitude intervals, or other types of multiplexed codes ([Akam and Kullmann, 2014](#); [Cover and Thomas, 2006](#); [Eyherabide and Samengo, 2010](#); [Gross et al., 2013](#); [Harvey et al., 2013](#); [Huk, 2012](#); [Oppenheim et al., 1997](#); [Panzeri et al., 2010](#); [Shannon and Weaver, 1949](#)). Moreover, neural activity may be characterized using continuous (quantitative) variables, as opposed to discrete (qualitative, nominal or categorical) variables. Continuous variables naturally arise when estimating firing rates, peak amplitudes, phases, or mean power within specific frequency-bands and time-intervals using tuning curves, local-field potentials (LFP), event-related potentials (ERP), or sensor signals from brain-imaging devices including electroencephalography (EEG) and magnetoencephalography (MEG). Unless quantized (discretized), these estimations may yield locally-smooth probability densities for which the above paradox need not arise.

To test this, we build another hypothetical experiment by replacing the discrete responses in [Fig. 3\(a\)](#) with continuous ones generated using unit-variance Gaussian distributions ([Fig. 4\(a\)](#)). Mathematically, their probability distributions can be rewritten as follows

$$P(\mathbf{R}_1|S_1) \propto \exp \left[\frac{(R_1^1 - R_1^2)^2}{2(\rho^2 - 1)} - \frac{(R_1^1 - \mu_1)(R_1^2 - \mu_1)}{1 + \rho} \right] \quad (20)$$

$$P(\mathbf{R}_2|S_2) \propto \exp \left[-0.5 (R_2^1 - \mu_2)^2 \right]. \quad (21)$$

where ρ is the correlation coefficient between R_1^1 and R_1^2 ; $\mu_1=2$ if $S_1=\square$ and 4 if $S_1=\bigcirc$; and $\mu_2=2$ if $S_2=A$ and 4 if $S_2=B$.

To study multiplexed codes, we reinterpret R_1^1 and R_1^2 as the amplitudes of 10 Hz -oscillations at two different cortical areas, denoted 1 and 2. Analogously, R_2^1 as the amplitude of 80 Hz -oscillations at the cortical area 1. Mathematically,

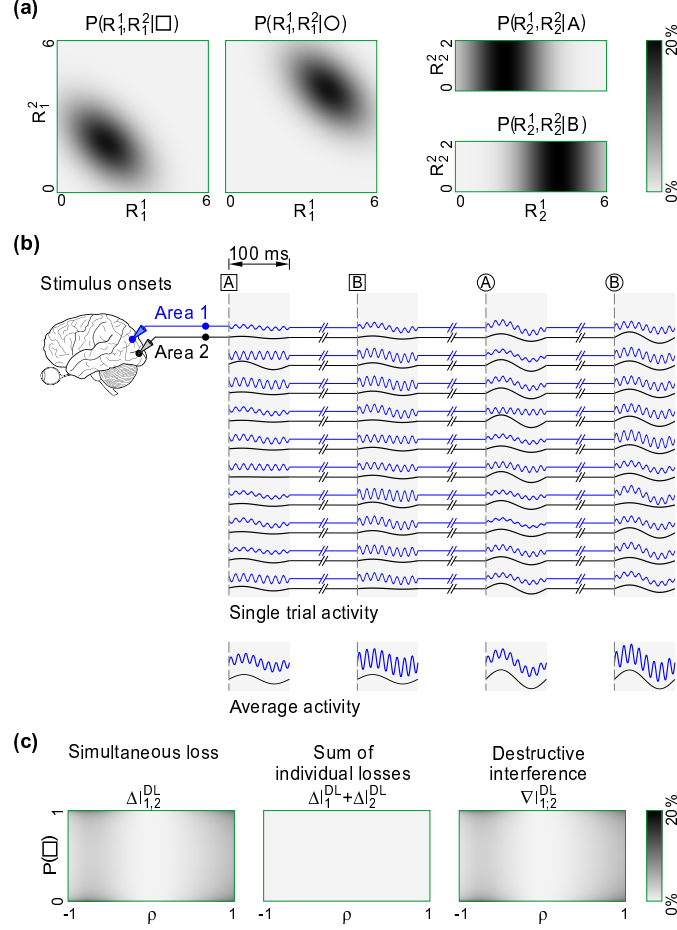


Figure 4: Destructive interference in brain models. (a) Cartesian representation of the probability densities defined in (20) and (21). The dummy variable R_2^2 uniformly distributed within the interval $[0, 2]$ was introduced for visualization purposes. (b) Hypothetical recording of the activity in two cortical areas, elicited within 100 ms after the onset of each stimulus \square , O , A and B , and defined through (22) and (23). (c) Analogous description to Fig. 3(c), but with $P([2, 2] | \square)$ replaced by ρ .

$$Z_1 = R_1^1 \sin(20 \pi t) + R_2^1 \sin(160 \pi t) \quad (22)$$

$$Z_2 = R_1^2 \sin(20 \pi t), \quad (23)$$

where Z_1 and Z_2 represent the possibly-filtered signals recorded from the cortical areas 1 and 2, and t is the time (Fig. 4(b)).

Analogous to the experiment in the previous section, noise correlations only occur between R_1^1 and R_1^2 . However, here they affect the amplitudes of the 10 Hz-oscillations at the two different cortical areas, as opposed to the spike counts of two neurons within the same population. Nevertheless, they need not be important from the decoding per-

spective regardless of their sign and strength, as previous studies have already shown (Averbeck et al., 2006; Averbeck and Lee, 2006; Eyherabide and Samengo, 2013).

Unfortunately, the results and conclusions of the above studies cannot be directly applied to our experiment because they were obtained using different measures and information notions. Nevertheless, we can rigorously prove that value of ΔI^{DL} computed taking only the 10 Hz-oscillations into account, here denoted ΔI_1^{DL} , is zero regardless of ρ . Indeed, after some algebra, $P(\mathbf{S}||[R_1^1, R_1^2])$ can be proved equal to $P^{NI}(\mathbf{S}||[R_1^1, R_1^2], \theta)$ for $\theta=1/(1+\rho)$ whenever $P([R_1^1, R_1^2])>0$.

According to previous studies, this result would indicate that noise correlations are irrelevant for decoding them in isolation regardless of the correlation strength. Even so, the value of $\Delta I_{1,2}^{DL}$ computed using all oscillations amplitudes in all cortical areas is always positive (except when $\rho=0$). Based on our results in the previous section, we can conclude that $\Delta I_{1,2}^{DL}$ is completely attributable to destructive interference, namely $\Delta I_{1,2}^{DL}=\nabla I_{1,2}^{DL}$ (Fig. 4(c)).

In conclusion, we have shown that destructive interference is ubiquitous regardless of the type of variables employed to characterize the recorded neural activity, and the neural substrates that the variables are interpreted to represent. In this way, we have conclusively answered two of the most recurrent questions in computational neuroscience: whether the phenomena observed when studying single neurons would also occur when studying cortical areas, and whether the results obtained with discrete variables would also emerge for continuous variables.

In addition, recall that we interpreted Fig. 4 as an example of multiplexed codes that transmit independent information through frequency division, whereas we can interpret Fig. 3 as an example of multiplexed codes that transmit independent information through space division (Oppenheim et al., 1997; Panzeri et al., 2010). These interpretations allow us to predict that destructive interference is a characteristic feature of multiplexed codes regardless of their implementation. However, our prediction need not imply that demultiplexing improves the performance of optimal NI decoders when information is independent, as we show in Section 3.9.

3.7 Ubiquity of destructive interference

In the previous section, we quantified ΔI^{DL} using only two populations transmitting independent information and found that it is superadditive. This result constitutes a major departure both from traditional views on the desirable properties of information measures, and from traditional expectations when operating on independent-information streams (Cover and Thomas, 2006; Fano, 1961; Oizumi et al., 2016; Schneidman et al., 2003). However, the aforementioned experiments may still seem overly simple, thereby

questioning the generality of our results. In this section, we trace back this phenomenon to the mathematical definition of ΔI^{DL} , as opposed to particular properties of the hypothetical neural data analyzed above.

Recall that ΔI^{DL} is defined through (7) as the minimization over the parameter θ of the function $\delta I(\theta)$, which is given by the following

$$\delta I(\mathbf{S}; \mathbf{R}|\theta) = \sum_{\mathbf{S}, \mathbf{R}} P(\mathbf{S}, \mathbf{R}) \ln \frac{P(\mathbf{S}|\mathbf{R})}{P^{NI}(\mathbf{S}|\mathbf{R}, \theta)}. \quad (24)$$

When two neural populations transmit independent information, like those in Fig. 3 and Fig. 4, it can be readily shown that, in addition to (9), the following equations hold

$$P(\mathbf{S}|\mathbf{R}) = P(S_1|\mathbf{R}_1) P(S_2|\mathbf{R}_2) \quad (25)$$

$$P^{NI}(\mathbf{S}|\mathbf{R}, \theta) = P^{NI}(S_1|\mathbf{R}_1, \theta) P^{NI}(S_2|\mathbf{R}_2, \theta), \quad (26)$$

where S_j denotes the stimulus features encoded in the activity \mathbf{R}_j of the j^{th} population. Therefore, we can rewrite (27) for these cases as follows

$$\delta I(\mathbf{S}; \mathbf{R}|\theta) = \delta I(\mathbf{S}_1; \mathbf{R}_1|\theta) + \delta I(\mathbf{S}_2; \mathbf{R}_2|\theta), \quad (27)$$

thereby proving that $\delta I(\theta)$ is additive when information is independent. For compactness, hereinafter we employ the abbreviated notation introduced in Section 2.7, and rewrite the above equation simply as $\delta I_{1,2}(\theta) = \delta I_1(\theta) + \delta I_2(\theta)$.

Notice that the additivity of $\delta I(\theta)$ directly implies the additivity of ΔI^D . Indeed ΔI^D can be computed as $\delta I(\theta)$ with $\theta=1$ (Latham and Nirenberg, 2005). Therefore, using the aforementioned notation, $\Delta I_{1,\dots,J}^D = \Delta I_{1+\dots+J}^D$. This result should not be confused with the additivity found in Nirenberg and Latham (2003), which involved neither decoders nor independent information, and is limited to non-overlapping response distributions.

However, the additivity of $\delta I(\theta)$ need not imply the additivity of ΔI^{DL} . Indeed, $\delta I(\theta)$ is convex (also called U-concave Gallager, 1968; Latham and Nirenberg, 2005), and the sum of the minima to two convex functions can never exceed the minimum of their sum. Therefore, the minimum of each term in (27), which correspond from left to right to $\Delta I_{1,2}^{DL}$, ΔI_1^{DL} and ΔI_2^{DL} , are related according to the following

$$\Delta I_{1,2}^{DL} \geq \Delta I_1^{DL} + \Delta I_2^{DL}. \quad (28)$$

Strict inequality holds whenever the minima occur at different locations, irrespective of their separation or the values of the minima (Fig. 5).

To illustrate this, we build the three examples shown in Fig. 5. The first example is based on Fig. 3(a) with $P(\square)$ and $P([2, 2]|\square)$ equal to 0.8 (Fig. 5(a)), and shows the

characteristic U-shape of $\delta I(\theta)$ regardless of whether it is computed using the neurons in population 1, in population 2, or in both. The minimum of $\delta I_1(\theta)$, namely ΔI_1^{DL} , occurs at $\theta \rightarrow \infty$, whereas the minimum of $\delta I_2(\theta)$, namely ΔI_2^{DL} , occurs at $\theta=1$. As expected, their sum is less than the minimum of $\delta I_{1,2}(\theta)$, namely $\Delta I_{1,2}^{DL}$, thereby leading to destructive interference.




Figure 5: Destructive interference stems from differences in the location of minima.

The second example is based on Fig. 4(a) with $P(\square)=0.8$ and $\rho=-0.5$ (Fig. 5(b)). This example is qualitatively analogous to the first one except that ΔI_1^{DL} occurs not at ∞ , but at $\theta=2$. The third example (Fig. 5(c)) is analogous to the second one, but with different correlation coefficients for \square (-0.8) and \bigcirc (-0.2). As a result, ΔI_1^{DL} is not zero, but positive and equal to $\approx 3.8\%$. In both cases, the observed differences are inessential and cannot prevent the occurrence of destructive interference.

Nevertheless, notice that a value of θ that simultaneously minimizes each and every term in (27) may theoretically occur. Indeed, recall the experiment in Fig. 3 with $P(\mathbf{S}, \mathbf{R})$ set as in Section 3.4. In that case, we found that $\Delta I_2^{DL}=0$ and $\Delta I_{1,2}^{DL}=\Delta I_1^{DL}$, from which we can immediately conclude that $\nabla I_{1,2}^{DL}=0$. This conclusion agrees with the fact that both $\delta I_1(\theta)$ and $\delta I_2(\theta)$ are both minimized for the same value of θ . The fact that the value of θ turns out to be unity is inessential.

Our results not only prove that the superadditivity of ΔI^{DL} stems directly from its mathematical definition, but also that the property is ubiquitous and independent of the type of data on which ΔI^{DL} is applied. Most importantly, they show that previous experimental findings in which ΔI^{DL} either grows with the number of neurons or with the decoding-window length or lies close to ΔI^D , need not be completely attributable to temporal correlations across time bins, pseudo-correlations caused by inappropriately assuming stationarity, or higher-order correlations, as previous studies have conjectured (Latham and Roudi, 2013; Oizumi et al., 2009, 2010). Instead, they can at least partially arise, even when information is independent, due to destructive interference.

Answering these questions requires that we test our results on arbitrary number J of neural populations. To that end, we use mathematical induction and rewrite $\Delta I_{1,\dots,J}^{DL}$ as the minimization of the convex function $\delta I_{1,\dots,J}(\theta)=\delta I_{1,\dots,J-1}(\theta)+\delta I_J(\theta)$, which minima are related through $\Delta I_{1,\dots,J}^{DL}\geq\Delta I_{1,\dots,J-1}^{DL}+\Delta I_J^{DL}$. This result implies that

$$\Delta I_{1,\dots,J}^{DL} \geq \Delta I_{1+\dots+J}^{DL}, \quad (29)$$

with equality if and only if some $\hat{\theta}$ exists that simultaneously minimizes $\delta I_j(\theta)$ for all $1 \leq j \leq J$ (that is, $\delta I_j(\hat{\theta}) = \min_{\theta} \delta I_j(\theta)$ for all $1 \leq j \leq J$).

Based on the above result, we can define destructive interference for an arbitrary number J of independent information streams as the following difference

$$\nabla I_{1,\dots,J}^{DL} = \Delta I_{1,\dots,J}^{DL} - \Delta I_{1+\dots+J}^{DL}. \quad (30)$$

It immediately follows from our demonstration of (29) that $\nabla I_{1,\dots,J}^{DL}$ never decreases with J .

In addition, the condition for equality in (29) immediately implies that the observation of $\nabla I_{1,\dots,J}^{DL} > 0$ requires of two independent information streams j_1 and j_2 , for which $\delta I_{j_1}(\theta)$ and $\delta I_{j_2}(\theta)$ are strictly convex, and achieve their minima at different values of θ . These conditions need not always hold, for $\delta I(\theta)$ can also be constant (Eyherabide and Samengo, 2013), as in the cases that fulfill the third property mentioned in Section 3.3. However, the fact that in these cases $\Delta I^{DL} = \Delta I^D$ need not imply that this condition should not hold for both streams. Indeed, destructive interference can arise even if only one stream exists for which $\Delta I^{DL} \neq \Delta I^D$, as shown in Fig. 5.

In conclusion, we have shown that destructive interference is a direct consequence of the convex minimization that defines the estimate ΔI^{DL} of the actual communication information loss ΔCI . In addition, we have derived the necessary and sufficient conditions for destructive interference to arise. These conditions need not always hold, as in Fig. 3 with the probabilities set as in Section 3.4, but our results show that the conditions are quite unrestrictive. Most importantly, we have extended the validity of our results to arbitrarily-complex independent information streams, regardless of their number and type.

3.8 Relative monotonic growth

We have just shown that $\nabla I_{1,\dots,J}^{DL}$ never decreases with the number J of independent information streams. However, this trend need not apply when measuring $\nabla I_{1,\dots,J}^{DL}$ relative to the transmitted information $I_{1,\dots,J}$ or to $\Delta I_{1,\dots,J}^{DL}$. Should it decrease instead, the destructive interference would become a minor component of ΔI^{DL} , and according to current beliefs, would play a minor role in the cost of ignoring noise correlations in optimal decoding. In this section, we test this hypothesis and show that the relative average of $\nabla I_{1,\dots,J}^{DL}$ never decreases, thereby driving $\Delta I_{1,\dots,J}^{DL}$ towards, but not necessarily reaching, its upper bound $\Delta I_{1,\dots,J}^D$.

To test this hypothesis, we build an experiment with J neural populations analogous to the first population in Fig. 3. These populations fire in response to visual stimuli $\mathbf{S}=[S_1, \dots, S_J]$ composed of J different and independently-chosen stimulus features. Each S_j denotes a frame (\square or \bigcirc) projected at a different location in a screen. For each value of S_j , the j^{th} population produces only two types of responses (Fig. 6).

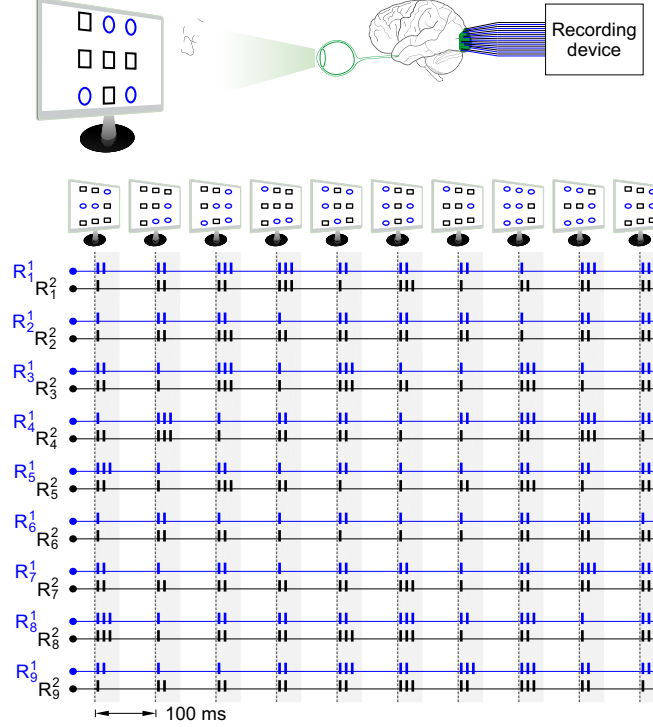


Figure 6: Hypothetical experiment with 9 populations of two neurons each firing independently and selectively to 9 different and independently chosen stimulus features. Each feature is a frame (\square or \bigcirc) located at a different position in the screen. The populations were generated by replicating the first population in Fig. 3(a), except that $P(S_j=\square)$ and $P(\mathbf{R}_j=[2, 2]|S_j=\square)$ were chosen independently for each population.

In this experiment, $P(\mathbf{S}, \mathbf{R})$ was chosen to fulfill (9). In this way, each j^{th} population fires independently of the other populations and selectively to the j^{th} stimulus feature S_j , thereby constituting independent information streams. For each stream, we set $P(\mathbf{R}_j=[2, 3]|S_j=\bigcirc)$ equal to $P(\mathbf{R}_j=[2, 2]|S_j=\square)$, and chose $P(S_j=\square)$ and $P(\mathbf{R}_j=[2, 2]|S_j=\square)$ both uniformly from the interval $[0.05, 0.95]$, and independently for each j^{th} population.

Repeating the above procedure, we generated 128 random instances of the same hypothetical experiment, each one with 1024 independent information streams. For each instance, we computed $\nabla I_{1;\dots;J}^{DL}$ using the first J streams. This computation showed that the relative value of $\nabla I_{1;\dots;J}^{DL}$ can increase or decrease with J depending on the

sample (Fig. 7(a), left). However, we also found that the average value of $\nabla I_{1;\dots;J}^{DL}$ across instances never decreased.

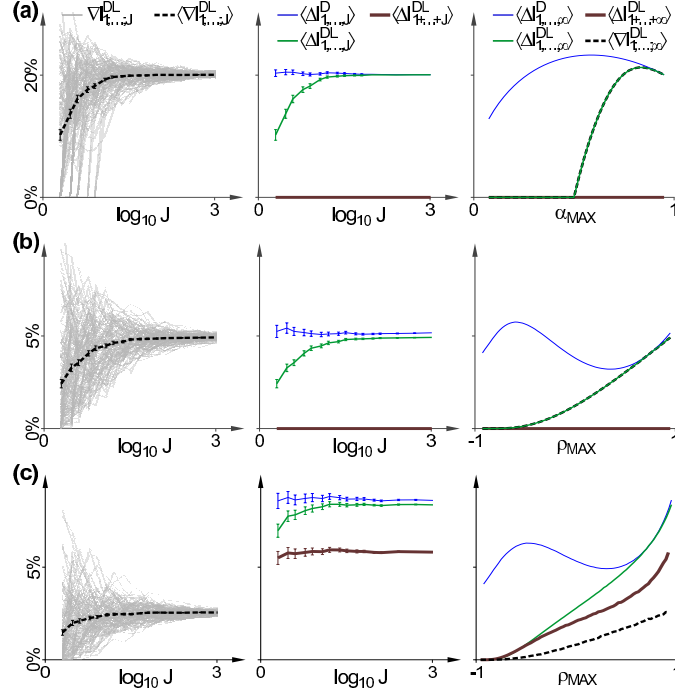


Figure 7: Destructive interference grows with the number of independent-information streams. (a) Variation as the number J of streams grows, for 128 instances of hypothetical experiments generated from Fig. 5(a), of the destructive interference (left), its average across instances (left) and the averages of the information losses $\langle \Delta I_{1,\dots,J}^D \rangle$, $\langle \Delta I_{1+,\dots,J}^{DL} \rangle$ and $\langle \Delta I_{1,\dots,J}^{DL} \rangle$ (middle). The right panel shows the value attained by the destructive interference and the above information losses in the limit for large J , when all $P(\mathbf{R}_j=[2,2]|S_j=\square)$ are drawn uniformly from the interval $[0.05, \alpha_{MAX}]$, as a function of α_{MAX} (right panel). All values are given relative to $I_{1,\dots,J}$. (b) Analogous description to (a), but for 128 instances of hypothetical experiments generated from Fig. 5(b), and with correlation coefficients in the right panel drawn from the interval $[-0.95, \rho_{MAX}]$. (c) Analogous description to (b), but for 128 instances of hypothetical experiments generated from Fig. 5(c). Vertical segments: standard error of the mean.

For the reasons we mentioned in Section 3.6, we repeated the analysis using as prototype information streams not the first population in Fig. 3, but the amplitudes of the 10 Hz -oscillations in the activity of two cortical areas studied in Fig. 4. The resulting experiments can be thought as concerning the information transmitted independently by the amplitudes of 10 Hz -oscillations at J pairs of cortical areas, namely

$$Z_j^1 = R_j^1 \sin(20 \pi t) \quad (31)$$

$$Z_j^2 = R_j^2 \sin(20 \pi t), \quad (32)$$

where Z_j^k denotes the activity recorded in the k^{th} cortical area of the j^{th} pair. However, the experiments can also be interpreted as the amplitudes of oscillations within J different frequency bands f_1, \dots, f_J in the activity of two cortical areas, namely

$$Z_1 = \sum_{j=1}^J R_j^1 \sin(2 \pi f_j t) \quad (33)$$

$$Z_2 = \sum_{j=1}^J R_j^2 \sin(2 \pi f_j t), \quad (34)$$

or a combination of both.

As before, the stimulus-response probabilities were chosen to fulfill (9). Each $P(\mathbf{R}_j|S_j)$ was given by the following unit-variance Gaussian distribution

$$P(\mathbf{R}_j|S_j) \propto \exp \left[\frac{(R_j^1 - R_j^2)^2}{2(\rho_j(S_j)^2 - 1)} - \frac{(R_j^1 - \mu_j)(R_j^2 - \mu_j)}{1 + \rho_j(S_j)} \right], \quad (35)$$

where $\rho_j(S_j)$ is the correlation coefficient between R_j^1 and R_j^2 given S_j ; whereas $\mu_j=2$ if $S_j=\square$ and 4 if $S_j=\bigcirc$. We chose $P(S_j=\square)$ and all correlation coefficients independently for each stream, the former uniformly from the interval $[0.05, 0.95]$, and the latter uniformly from the interval $[-0.95, 0.95]$. Analogously to Fig. 5, we considered both the case in which the correlation coefficients $\rho_j(R_j^1, R_j^2|\square)$ and $\rho_j(R_j^1, R_j^2|\bigcirc)$ for each j^{th} information stream coincide (Fig. 7(b)), and the case in which $\rho_j(R_j^1, R_j^2|\square)$ and $\rho_j(R_j^1, R_j^2|\bigcirc)$ differ (Fig. 7(c)). In both cases, we found trends for the relative value and the relative average of $\nabla I_{1,\dots,J}^{DL}$ as a function of J analogous to those in Fig. 7(a).

We also found that, for the three hypothetical experiments mentioned above, the relative averages of ΔI^D and $\Delta I_{1+\dots+J}^{DL}$ remain virtually constant as J grows (Fig. 7, middle panels). In that case, the growth with J of the average $\nabla I_{1,\dots,J}^{DL}$ relative to $I_{1,\dots,J}$ is equivalent to that of the average $\Delta I_{1,\dots,J}^{DL}$ relative to $I_{1,\dots,J}$. Most importantly, these results imply that the relative value of $\nabla I_{1,\dots,J}^{DL}$ with respect to $\Delta I_{1,\dots,J}^{DL}$ also grows with J . In other words, the proportion of $\Delta I_{1,\dots,J}^{DL}$, as opposed to that of $I_{1,\dots,J}$, explained by $\nabla I_{1,\dots,J}^{DL}$, never decreases with the number of independent information streams.

The above results are valid not only for the three hypothetical experiments studied in Fig. 7, but for more general hypothetical experiments comprising J independent information streams defined through independent and identically distributed parameter vectors. For these class of experiments, we can prove that the following relations hold

$$\langle I_{1,\dots,\lambda J} \rangle = \lambda J \langle I_j \rangle \quad (36)$$

$$\langle \Delta I_{1,\dots,\lambda J}^D \rangle = \lambda J \langle \Delta I_j^D \rangle \quad (37)$$

$$\langle \Delta I_{1+\dots+\lambda J}^{DL} \rangle = \lambda J \langle \Delta I_j^{DL} \rangle \quad (38)$$

$$\langle \Delta I_{1,\dots,\lambda J}^{DL} \rangle \geq \lambda \langle \Delta I_{1,\dots,J}^{DL} \rangle \quad (39)$$

$$\langle \nabla I_{1,\dots;\lambda J}^{DL} \rangle \geq \lambda \langle \nabla I_{1,\dots;J}^{DL} \rangle, \quad (40)$$

where λ is a positive integer. All these equations rest partially on the linearity of the mean. In addition, (36) stems from (10); (37), from (27) with $\theta=1$; (39), from the convexity of δI ; and (40), from subtracting (39) and (38).

Dividing by $\langle I_{1,\dots,\lambda J} \rangle$ each side of (37)–(40) yields the following relations for the relative averages

$$\langle \Delta I_{1,\dots,\lambda J}^D \rangle / \langle I_{1,\dots,\lambda J} \rangle = \langle \Delta I_j^D \rangle / \langle I_j \rangle \quad (41)$$

$$\langle \Delta I_{1+\dots+\lambda J}^{DL} \rangle / \langle I_{1,\dots,\lambda J} \rangle = \langle \Delta I_j^{DL} \rangle / \langle I_j \rangle \quad (42)$$

$$\langle \Delta I_{1,\dots,\lambda J}^{DL} \rangle / \langle I_{1,\dots,\lambda J} \rangle \geq \langle \Delta I_{1,\dots,J}^{DL} \rangle / \langle I_{1,\dots,J} \rangle \quad (43)$$

$$\langle \nabla I_{1,\dots;\lambda J}^{DL} \rangle / \langle I_{1,\dots,\lambda J} \rangle \geq \langle \nabla I_{1,\dots;J}^{DL} \rangle / \langle I_{1,\dots,J} \rangle. \quad (44)$$

Furthermore, dividing (40) by $\langle \Delta I_{1,\dots,\lambda J}^{DL} \rangle$ yields the following relation

$$\langle \nabla I_{1,\dots;\lambda_1 \lambda_2 J}^{DL} \rangle / \langle \Delta I_{1,\dots,\lambda_1 \lambda_2 J}^{DL} \rangle \geq \langle \nabla I_{1,\dots;\lambda_1 J}^{DL} \rangle / \langle \Delta I_{1,\dots,\lambda_2 J}^{DL} \rangle, \quad (45)$$

after replacing λ with the product of two positive integers λ_1 and λ_2 , and some relatively simple algebra. These relations generalize our observations in Fig. 7, and prove that the destructive interference never decreases in absolute or relative magnitude, and becomes increasingly important as the number of independent information streams grows.

We also noticed that the relative magnitude of $\nabla I_{1,\dots;J}^{DL}$ converges as J grows to a value, here called $\nabla I_{1,\dots;\infty}^{DL}$. This observation need not be surprising for the relative average of $\nabla I_{1,\dots;\infty}^{DL}$ across instances. Indeed, not only this value is bounded by unity, but also increasing with J , as we have shown, and therefore the monotone convergence theorem ensures that a limit exists as $J \rightarrow \infty$. However, our observation also applies to each instance of the experiments, as opposed to their averages (Fig. 7, left). Most importantly, these observations open up the possibility that $\nabla I_{1,\dots;J}^{DL}$ becomes so large when $J \rightarrow \infty$ as to drive $\Delta I_{1,\dots,J}^{DL}$ close to its maximum value, namely $\Delta I_{1,\dots,J}^D$.

To test this hypothesis, we generated different instances of the experiment in Fig. 7(a) by choosing $P(\mathbf{R}_j=[2, 2] | S_j=\square)$ uniformly from intervals $[0.05, \alpha_{MAX}]$, for $0.05 \leq \alpha_{MAX} \leq 0.95$. For these instances, we computed $I(\cdot, 1, \dots, \infty)$ analytically, and estimated $\Delta I_{1,\dots,\infty}^D$, $\Delta I_{1,\dots,\infty}^{DL}$ and $\Delta I_{1+\dots+\infty}^{DL}$ numerically using Matlab R2015b.

We found that the relative value of $\nabla I_{1,\dots,\infty}^{DL}$, and consequently of $\Delta I_{1,\dots,\infty}^{DL}$, reached its maximum when $\alpha_{MAX} \approx 0.83$. This value was estimated as the maximum of a cubic function that locally approximated the trace of $\nabla I_{1,\dots,\infty}$ as α_{MAX} grows. Furthermore, we found that both $\nabla I_{1,\dots,\infty}^{DL}$ and $\Delta I_{1,\dots,\infty}^{DL}$ converged to $\Delta I_{1,\dots,\infty}^D$ whenever $P(\mathbf{R}_j=[2, 2]|S_j=\square)$ was drawn from intervals centered at 0.5. Indeed, for large J , $\delta I_{1,\dots,J}(\theta)$ can be approximated using the law of large numbers as follows

$$\delta I_{1,\dots,J}(\theta) \approx J \left\langle q_j \alpha_j \ln \left(1 + \frac{\bar{q}_j}{q_j} \left(\frac{\bar{\alpha}_j}{\alpha_j} \right)^\theta \right) \right\rangle, \quad (46)$$

where for compactness, we have employed a notation analogous to that introduced after (18), defining $q_j = P(S_j=\square)$, $\alpha_j = P(\mathbf{R}_j=[2, 2]|S_j=\square)$, and $\bar{x} = 1 - x$ for any real value x . After some algebra, this equation can be shown to reach its minimum when $\theta = 1$ provided that q_j and α_j are chosen with probability distributions that are symmetric about 0.5, and hence

$$\Delta I_{1,\dots,J}^{DL} \underset{J \gg 1}{\approx} \Delta I_{1,\dots,J}^D. \quad (47)$$

In Fig. 7(b)-(c), different instances were generated by choosing the correlation coefficients uniformly from intervals $[-0.95, \rho_{MAX}]$, for $-0.95 \leq \rho_{MAX} \leq 0.95$. In these cases, we estimated $I_{1,\dots,\infty}$, $\Delta I_{1,\dots,\infty}^D$ and $\Delta I_{1,\dots,\infty}^{DL}$ using Monte Carlo integration in Python 3.4.3 with the packages Vegas 3.0 and SciPy 0.14.1, using 2000000 vector samples, divided in 10 iterations for training and 10 iterations for evaluation. The value of $\Delta I_{1,\dots,\infty}^{DL}$ was estimated using 163840 independently generated streams.

Unlike in Fig. 7(a), here the relative values of both $\nabla I_{1,\dots,\infty}^{DL}$ and $\Delta I_{1,\dots,\infty}^{DL}$ increased with ρ_{MAX} . However, $\Delta I_{1,\dots,\infty}^{DL}$ converged to $\Delta I_{1,\dots,\infty}^D$ only when $\rho_{MAX} \approx 0.73$ in both Fig. 7(b) and Fig. 7(c). This value was estimated analogously to the value of α_{MAX} in Fig. 7(a).

To summarize, our results show that as the number of independent information streams grows, the relative amount of destructive interference is never negative and never decreases. This result finally answers the question we posited in the previous section, and shows that ΔI^{DL} may grow with the number of neurons or the decoding-window duration by the virtue of destructive interference alone, or in conjunction with other possible causes. In addition, destructive interference may drive the loss $\Delta I_{1,\dots,J}^{DL}$ near its upper bound $\Delta I_{1,\dots,J}^D$. Analogous observations have previously been regarded as a sign that $\Delta I_{1,\dots,J}^D$ is tight (Latham and Roudi, 2013). Our results once again cast doubt on this interpretation, but this time by showing that it may stem from a paradoxical growth in the cost of ignoring noise correlations.

3.9 The construction of optimal NI decoders is not unique

The measure $\Delta I_{1,\dots,J}^{DL}$ computed above stems from computing the measure ΔI^{DL} introduced by [Latham and Nirenberg \(2005\)](#) using all neurons in J populations that transmit independent information. According to previous studies, this computation would yield the exact information loss caused when optimally decoding the aforementioned neurons ignoring noise correlations. However, in this section we show that ΔI^{DL} is actually limited to a specific construction of optimal NI decoders, thereby opening up the possibility that other constructions be more efficient than predicted by ΔI^{DL} .

Specifically, we noticed that ΔI^{DL} , and consequently $\Delta I_{1,\dots,J}^{DL}$, use what we here call joint NI decoders (also called centralized integration; [Zhang et al., 2016](#)). These decoders produce simultaneous estimates \mathbf{S}^J of all stimulus features only after reading the concurrent responses \mathbf{R} of all populations ([Fig. 8\(a\)](#)). Mathematically, joint NI decoders can be defined as follows

$$\mathbf{S}^J = \arg \max_{\mathbf{S}} P^{NI}(\mathbf{S}|\mathbf{R}). \quad (48)$$

However, joint NI decoders need not be the only way in which optimal NI decoders can be constructed.

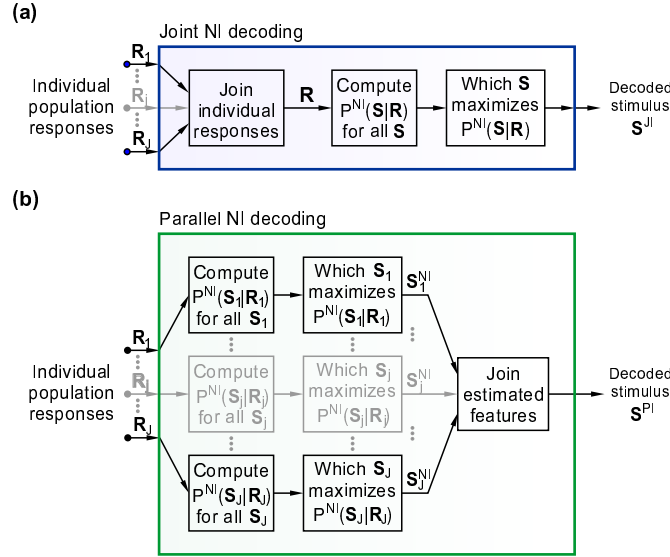


Figure 8: Two different constructions of optimal NI decoders. (a) Joint NI decoders choose simultaneously all the stimulus features that maximize the NI posterior $P^{NI}(\mathbf{S}|\mathbf{R})$. (b) Parallel NI decoders choose each stimulus feature singly but in parallel by maximizing each of their corresponding NI posteriors $P^{NI}(S_j|\mathbf{R}_j)$.

To show this, recall our definition of independent information given in [Section 2.7](#). According to this definition, J neural populations transmit independent information

when they fire independently and selectively to J independent stimulus features, respectively. In that case, the stimulus features can also be optimally identified in parallel (also known as distributed, modular or decentralized decoding; [Zhang et al., 2016](#)) regardless of whether neurons are noise independent or not. However, the use of parallel decoders when studying the role of noise correlations in neural decoding has previously been controversial ([Latham and Nirenberg, 2005](#); [Meister and Hosoya, 2001](#); [Nirenberg et al., 2001a](#); [Schneidman et al., 2003](#)).

As we note here, this controversy could have been avoided should previous studies have combined the outputs of the parallel decoders into a single consistent estimate. However, such combination may not be possible without taking implicitly or explicitly correlations into account or making additional assumptions ([Eyherabide and Samengo, 2013](#); [Jaynes, 2003](#); [Knill and Richards, 1996](#); [Landy et al., 1995](#); [Schneidman et al., 2003](#)). These difficulties do not arise here because populations transmit independent information.

Specifically, here we construct parallel NI decoders as two-stage processes ([Fig. 8\(b\)](#)). The first stage consists of J optimal NI decoders, each of which operates separately on a different population. Mathematically, the j^{th} optimal NI decoder reads the response \mathbf{R}_j of the j^{th} population and produces an estimate S_j^{PI} of the j^{th} feature, according to the following

$$S_j^{PI} = \arg \max_{S_j} P^{NI}(S_j | \mathbf{R}_j) = \arg \max_{S_j} P^{NI}(S_j | \mathbf{R}). \quad (49)$$

These estimates are concurrently fed into the second stage, which concatenates them to produce the estimated stimulus $\mathbf{S}^{PI} = [S_1^{PI}, \dots, S_J^{PI}]$. Unlike [Zhang et al. \(2016\)](#), the definitions of joint and parallel decoders given here are both valid for arbitrary stimulus distributions.

The last equality in (49) shows that feeding each j^{th} optimal NI decoder in the first stage only with the responses of the j^{th} population (or demultiplexing, as mentioned in [Section 3.6](#)) is inessential. Indeed, after some algebra, (14) turns into $P^{NI}(S_j | \mathbf{R}_j) = P^{NI}(S_j | \mathbf{R})$ for populations that transmit independent information. Therefore, parallel estimations of each individual stimulus feature S_j^{PI} can be conducted without the interference of other information streams that seemingly affect the joint estimations, at least according to the results we have obtained above.

To summarize, we have shown for the first time that ΔI^{DL} is only based on joint NI decoders. However, even though unnoticed by previous studies, we found that the construction of optimal NI decoders need not be unique, and different alternatives may exist depending on the case. In particular, we showed that, when information is independent, optimal NI decoding can be performed in parallel. These possibilities have remained largely ignored in the neural coding literature, and most importantly, may lead to con-

structions that outperform joint NI decoders and overcome the destructive interference. In that case, ΔI^{DL} would potentially overestimate ΔCI , and the importance of noise correlations in optimal decoding, for yet another reason than those we have mentioned above.

3.10 Decoding more efficiently than predicted

Our finding that the construction of optimal NI decoders is not uniquely defined immediately raises the question of what difference do different constructions make. To address this question, we compute in this section the information losses caused by parallel NI decoders using the information notions underlying ΔI^D , ΔI^B and ΔI^{DL} . Contrary to current pervasive ideas, these computations need not stem straight-forwardly from previous results on encoded information losses for the reasons we mentioned in [Section 3.1](#).

The axiomatic information loss caused by parallel NI decoders can be computed by reinterpreting the first stage as parallel transformations of each individual population response. When information is independent, [\(14\)](#) holds, and the axiomatic information losses $\Delta I_1^{DL}, \dots, \Delta I_J^{DL}$ caused by each of these parallel transformations are additive. The second stage is invertible, and hence lossless. Consequently, using the notation introduced in [Section 2.7](#), the axiomatic information loss caused by parallel NI decoders is equal to $\Delta I_{1+\dots+J}^B$.

Analogous results can be obtained for the value of ΔI^D associated with parallel NI decoders. To that end, recall that ΔI^D was derived by [Nirenberg et al. \(2001a\)](#) using a notion of information with roots in coding theory, here called descriptive information. Within this notion, ΔI^D can be interpreted as the increment caused by ignoring noise correlations in the minimum average description length L_S of the stimulus identity S after observing the population response R ([Cover and Thomas, 2006](#)).

Within that context, parallel NI decoding can be interpreted as describing S by concatenating the individual descriptions of its J stimulus features. The resulting minimum average description length L_S^{PI} equals the sum of the minimum average description lengths L_{S_1}, \dots, L_{S_J} associated with each stimulus feature, and therefore, their corresponding increments $\Delta I_1^D, \dots, \Delta I_J^D$ caused by ignoring noise correlations are additive. Consequently, the descriptive information loss caused by parallel NI decoders equals $\Delta I_{1+\dots+J}^D$. Notice that this result need not immediately arise from the additivity of ΔI^D found in [Section 3.7](#), for that mathematical findings need not be conceptually related to the operation of parallel NI decoders.

The above results render seemingly reasonable to hypothesize that the communication information loss produced by parallel NI decoders equals $\Delta I_{1+\dots+J}^{DL}$. To rigor-

ously prove this, recall the notation introduced in [Section 2.7](#), and the derivation of [Latham and Nirenberg \(2005\)](#) introduced in [Section 3.1](#). According to their derivation, we can associate each j^{th} population with codebooks of up to $\exp(N \tilde{I}_j)$ sequences $\mathbf{S}_j^{[N]}$, where $\tilde{I}_j \leq I_j - \Delta I_j^{DL}$, for which the average decoding-error rate $P_{e,j}^{[N]}$ produced by optimal NI decoders vanishes exponentially as N grows.

Combining the above codebooks using Cartesian products yields a product codebook of up to $\exp(N \sum_{j=1}^J \tilde{I}_j)$ sequences. For these codebooks, Boole's inequality ([Casella and Berger, 2002](#)) yields after some algebra that the average decoding error rate $P_e^{[N]}$ decays exponentially as N grows at least as fast as $\max_j P_{e,j}^{[N]}$. Hence, $\Delta I_{1+\dots+J}^{DL}$ is achievable ([Cover and Thomas, 2006](#)) and quantifies the communication information loss caused when decoding J independent information streams ignoring noise correlations.

In conclusion, we have proved that, when information is independent, the information loss caused by parallel NI decoders is equal to the sum of the information losses caused by each of its constituent optimal NI decoders regardless of the underlying information notion. Most importantly, this result shows that $\Delta I_{1+\dots+J}^{DL}$ is achievable, thereby proving that parallel NI decoders can overcome the destructive interference. Therefore, we conclude that parallel NI decoders can be more efficient than predicted by $\Delta I_{1,\dots,J}^{DL}$, and that, contrary to current beliefs, noise correlations can be less important than predicted by ΔI^{DL} .

3.11 Joint NI decoders can potentially outperform parallel NI decoders

We have shown in the previous sections have shown that the communication information loss $\Delta I_{1+\dots+J}^{DL}$ caused by parallel NI decoders is never greater and can be less than the communication information loss $\Delta I_{1,\dots,J}^{DL}$ caused by joint NI decoders. These results seemingly indicate both that parallel NI decoders do outperform joint NI decoders, at least in terms of communication information losses, and that, regardless of how paradoxical it may seem, destructive interference need not merely stems from flaws in ΔI^{DL} . In this section, we address these hypotheses and show that joint NI decoders can both achieve $\Delta I_{1+\dots+J}^{DL}$ and potentially outperform parallel NI decoders, at least in terms of axiomatic information loss.

To prove that joint NI decoders can achieve $\Delta I_{1+\dots+J}^{DL}$, recall that the derivation of ΔI^{DL} in [Latham and Nirenberg \(2005\)](#) is based on the average decoding error probability $P_e^{[N]}$, which they wrote as $P \left(P^{NI}(\mathbf{S}|\mathbf{R}) < P^{NI}(\tilde{\mathbf{S}}|\mathbf{R}) \right)$. However, notice that the estimates \mathbf{S}^{JI} and \mathbf{S}^{PI} produced by joint and parallel NI decoders, respectively, are always associated with the same NI posteriors, namely

$$\begin{aligned}
P^{NI}(\mathbf{S}^J|\mathbf{R}) &= \max_{S_1, \dots, S_J} \prod_{j=1}^J P^{NI}(S_j|\mathbf{R}_j) \\
&= \prod_{j=1}^J \max_{S_j} P^{NI}(S_j|\mathbf{R}_j) \\
&= P^{NI}(\mathbf{S}^{PI}|\mathbf{R}).
\end{aligned} \tag{50}$$

Therefore, $P(P^{NI}(\mathbf{S}|\mathbf{R}) < P^{NI}(\mathbf{S}^J|\mathbf{R})) = P(P^{NI}(\mathbf{S}|\mathbf{R}) < P^{NI}(\mathbf{S}^{PI}|\mathbf{R}))$, and consequently joint and parallel NI decoders produce the same average decoding error probability for the same set of codebooks. In other words, joint NI decoders can also achieve $\Delta I_{1+\dots+J}^{DL}$, thereby rendering $\Delta I_{1,\dots,J}^{DL}$ as an upper bound of the actual communication information loss caused by joint NI decoders.

The equality between $P^{NI}(\mathbf{S}^J|\mathbf{R})$ and $P^{NI}(\mathbf{S}^{PI}|\mathbf{R})$ also implies that both decoders will typically produce the same estimates. As a result, the axiomatic information losses $\Delta I_{1,\dots,J}^B$ and $\Delta I_{1+\dots+J}^B$ caused by joint and parallel NI decoders, respectively, will typically coincide. However, differences in their estimates, and hence in their axiomatic information losses, can still stem from arbitrary tie-breaking rules.

Specifically, the maxima in (48) and (49), respectively, oftentimes occur for multiple stimuli. Hence, joint and parallel NI decoders cannot unambiguously produce an estimate. In practice, these situations need not be rare and may arise, for example, due to probability quantization when estimating them from experimental frequencies (Casella and Berger, 2002; Samengo, 2002). Should ambiguities arise, they can be resolved by adopting tie-breaking rules. Depending on how they are set, $\Delta I_{1,\dots,J}^B$ can be greater, equal or less than $\Delta I_{1+\dots+J}^B$.

To illustrate this, we build the hypothetical experiment shown in Fig. 9(a). Analogously to Fig. 3(a), this experiment consists of two populations that fire independently and selectively to independent stimulus features. However, here population 1 also produces the response $\mathbf{R}_1=[2, 2]$ after \bigcirc (where $\mathbf{R}_j=[R_j^1, R_j^2]$). In addition, population 2 includes a second neuron that produces the same number of spikes as the first neuron after A, and only two spikes regardless of the response of the first neuron after B (Fig. 9(b)).

Suppose that we set $P(S_1, \mathbf{R}_1)$ so that $P(\square, [2, 2])=12/30$, $P(\bigcirc, [2, 2])=7/30$ and $P(\bigcirc, [2, 3])=2/30$. In that case, $P^{NI}(\square|\mathbf{R}_1)=P^{NI}(\bigcirc|\mathbf{R}_1)$ only when $\mathbf{R}_1=[2, 2]$, and thus optimal NI decoders cannot unambiguously choose a frame. To resolve this ambiguity, we can adopt the convention of choosing \bigcirc over \square , here denoted $\bigcirc > \square$. This convention minimizes the axiomatic information loss ΔI_1^B computed using only population 1 ($\approx 31\%$).

Analogously, suppose that we set $P(S_2, \mathbf{R}_2)$ so that $P(A, [2, 2])=90/190$ and $P(B, [2, 2])=81/190$. In that case, $P^{NI}(A|\mathbf{R}_2)=P^{NI}(B|\mathbf{R}_2)$ only when $\mathbf{R}_2=[2, 2]$, and

thus optimal NI decoders cannot unambiguously choose a letter. To resolve this ambiguity, we can adopt the convention $A > B$, which minimizes the axiomatic information loss ΔI_2^B computed using only population 2 ($\approx 11\%$).

Together, the above two conventions minimize the axiomatic information loss ΔI_{1+2}^B caused by parallel NI decoders ($\approx 42\%$). On the contrary, the opposite conventions (i.e., $\bigcirc < \square$ and $A < B$) maximize ΔI_{1+2}^B ($\approx 46\%$; with $\Delta I_1^B \approx 34\%$ and $\Delta I_2^B \approx 12\%$). Other two ways of combining the conventions exist that yield intermediate values of ΔI_{1+2}^B , thereby adding up to four different constructions of parallel NI decoders.




Figure 9: Experiment in which joint and parallel NI decoders can yield different information losses even though information is independent. (a) Analogous description to Fig. 3(a). (b) Analogous description of Fig. 3(b).

On the contrary, the number of joint NI decoders that can be constructed by choosing different tie-breaking rules is 64. These rules must choose between

1. \square and \bigcirc if $\mathbf{R}=[2, 2, 1, 1]$, because $P^{NI}(\square|\mathbf{R})=P^{NI}(\bigcirc|\mathbf{R})=0.5$;
2. \square and \square if $\mathbf{R}=[1, 1, 2, 2]$, because $P^{NI}(\square|\mathbf{R})=P^{NI}(\square|\mathbf{R})=0.5$;
3. \bigcirc and \bigcirc if $\mathbf{R}=[2, 3, 2, 2]$ or $[3, 2, 2, 2]$, because $P^{NI}(\bigcirc|\mathbf{R})=P^{NI}(\bigcirc|\mathbf{R})=0.5$;
4. \square and \bigcirc if $\mathbf{R}=[2, 2, 2, 3]$, because $P^{NI}(\square|\mathbf{R})=P^{NI}(\bigcirc|\mathbf{R})=0.5$; and
5. all stimuli if $\mathbf{R}=[2, 2, 2, 2]$, because $P^{NI}(\mathbf{S}|\mathbf{R})=0.25$ for all \mathbf{S} .

By adopting tie-breaking rules independently for each of the above five situations, we can find joint NI decoders that cause axiomatic information losses $\Delta I_{1,2}^B$ as low as $\approx 16\%$ (only $\approx 38\%$ of the minimum ΔI_{1+2}^B) or as large as $\approx 58\%$ (and thus $\approx 28\%$ larger than the maximum ΔI_{1+2}^B).

In conclusion, we have shown that joint and parallel NI decoders are almost always equivalent, not only in terms of axiomatic information losses, namely

$\Delta I_{1,\dots,J}^B = \Delta I_{1+\dots+J}^B$, but also in their stimulus estimates, namely \mathbf{S}^J and \mathbf{S}^{PI} . However, this need not be the case when tie-breaking rules must be chosen. In those cases, $\Delta I_{1,\dots,J}^B$ can be greater or less than $\Delta I_{1+\dots+J}^B$ depending on the chosen conventions, but, without restrictions, the best parallel NI decoder can never outperform the best joint NI decoder.

These results may also apply to communication information losses $\Delta I_{1+\dots+J}^{DL}$ and $\Delta I_{1,\dots,J}^{DL}$. Indeed, we have shown that both joint and parallel NI decodes can achieve the same communication information losses, namely $\Delta I_{1+\dots+J}^{DL}$. However, this need not always be the case for at least three reasons. First, neither our computations nor the derivation of ΔI^{DL} in [Latham and Nirenberg \(2005\)](#) have taken into account the potential effect of tie-breaking rules. Second, here we showed that the set of all codebooks employed by [Latham and Nirenberg \(2005\)](#) to derive ΔI^{DL} may produce larger average decoding error probabilities than some of its subsets. Third, the set of codebooks for which joint and parallel NI decoders achieve the minimum average decoding error probability, respectively, need not coincide.

These three considerations can be proved unnecessary when studying other response aspects such as spike counts or latencies. However, for the reasons we mentioned in [Section 3.1](#), this observation need not apply to the study of noise correlations from the decoding perspective. Consequently, our results open up the possibility that joint NI decoders outperform parallel NI decoders in terms of communication information loss. Most importantly, they allow us to finally refine the estimation of ΔCI as follows

$$\Delta CI \begin{cases} = 0 & \text{if } P(\mathbf{S} \neq \tilde{\mathbf{S}}) = 0 \\ \leq \Delta I_{1+\dots+J}^{DL} & \text{if information is independent} \\ \leq \Delta I^{DL} & \text{otherwise} \end{cases} \quad (51)$$

Above all, our results in this section show that the overestimation of ΔI^{DL} need not be constrained to independent information streams, and that current beliefs in the exactness of ΔI^{DL} may lead one to overestimate the importance of noise correlations in optimal decoding.

4 Discussion

Many measures have been proposed to quantify the information loss caused by ignoring noise correlations in neural decoding, but their conceptual and quantitative accuracy remains controversial ([Averbeck et al., 2006](#); [Eyherabide and Samengo, 2013](#); [Ince et al., 2010](#); [Latham and Nirenberg, 2005](#); [Latham and Roudi, 2013](#); [Meister and Hosoya, 2001](#); [Nirenberg et al., 2001a](#); [Nirenberg and Latham, 2003](#); [Oizumi et al., 2009, 2010](#);

Schneidman et al., 2003). Resolving these controversies is fundamental for understanding the role of noise correlations in brain computations over multiple groups of neurons or neural substrates, or when decoding brain signals recorded from multiple brain locations either in one or even multiple subjects (Babiloni and Astolfi, 2014; Hari and Kujala, 2009). In this study, we focus on one of the most prominent measures, due to its information-theoretical foundations and its underlying communication notion of information, here called ΔI^{DL} . This measure was introduced by Latham and Nirenberg (2005) based on the work of Merhav et al. (1994) on mismatched decoding, and it is currently considered the exact information loss caused by ignoring noise correlations in optimal decoding (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2016, 2009, 2010), despite the fact that, to our knowledge, the properties of ΔI^{DL} and the consequences of its putative exactness have both remained largely unexplored.

Accordingly, our first step in this direction was to address the implications of what to our knowledge is the only controversial finding to date concerning ΔI^{DL} and ΔI^B . As a putative exact measure of information loss, ΔI^{DL} would be expected to never exceed ΔI^B (Cover and Thomas, 2006; Quiroga and Panzeri, 2009; ?). However, as we have recently shown, this need not be the case (Eyherabide and Samengo, 2013). However, due to the rigorous derivation of ΔI^{DL} , it remains unclear whether this observation indicates major departures from traditional relations between information and decoding or flaws in ΔI^{DL} .

Although puzzling, here we argued that these observations need not actually contradict the belief that ΔI^{DL} is exact for at least two reasons. First, as mentioned in Section 3.1, ΔI^{DL} and ΔI^B are based on different notions of information, which differences are not necessarily unknown, but previous studies have often overlooked (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2010; Quiroga and Panzeri, 2009; Thomson and Kristan, 2005). Second, here we point out that even though the relation $\Delta I^{DL} < \Delta I^B$ holds when studying response aspects such as spike counts or latencies (Eyherabide, 2016), this observation need not immediately imply that the relation must also hold when studying noise correlations.

These two reasons led us to hypothesize that previous observations of ΔI^{DL} exceeding ΔI^B may simply stem from fundamental differences between the two underlying information notions rather than from previously unforeseen flaws in ΔI^{DL} . To disentangle these two possibilities, we compared for the first time the value of ΔI^{DL} with a direct computation of the communication information loss caused by ignoring noise correlations in optimal decoding that we called ΔCI . Contrary to currently thought (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi,

2013; Oizumi et al., 2016, 2009, 2010), our results showed for the first time that ΔI^{DL} need not be exact and can overestimate ΔCI at least when optimal NI decoders can perfectly identified some stimulus feature.

Using populations that transmit independent information, we also showed for the first time that, that ΔI^{DL} is actually superadditive. Specifically, we found that the value of ΔI^{DL} computed using all neurons in all populations is larger than the sum of the values of ΔI^{DL} computed using all neurons in each population, respectively. This result constitutes a major departure from the traditional additivity of mutual information (Cover and Thomas, 2006; Fano, 1961; Shannon and Weaver, 1949), thereby questioning current beliefs in the exactness of ΔI^{DL} (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2016, 2009, 2010).

This paradoxical increment in ΔI^{DL} was here called destructive interference, and shown ubiquitous regardless of whether the populations were interpreted as spatially or temporally multiplexed independent-information streams (Oppenheim et al., 1997; Panzeri et al., 2010), which we explained by tracing them back to the convex minimization in the definition of ΔI^{DL} (Eyherabide and Samengo, 2013; Latham and Nirenberg, 2005; Oizumi et al., 2016, 2009, 2010). Our explanation implies that may arise regardless of the correlation importance within populations, the composition of neural substrates and the type of multiplexed codes. Furthermore, the mathematical nature of this result extends this phenomenon beyond neural coding to information streams of arbitrary type (Ernst and Banks, 2002; Eyherabide and Samengo, 2010; Fano, 1961; Oizumi et al., 2016; Panzeri et al., 2010; Shannon and Weaver, 1949; Zhang et al., 2016).

We also found that the destructive interference grows with the number of populations, driving ΔI^{DL} towards ΔI^D (Fig. 7), and possibly reaching $\approx 100\%$ of the transmitted information (Fig. 3), even when each population could be safely decoded ignoring noise correlations (figures 3; 4; 5; 7). These results are qualitatively similar to previous experimental findings in which ΔI^{DL} grows with the number of neurons or with the decoding-window length, and lies close to ΔI^D (Latham and Roudi, 2013; Oizumi et al., 2009, 2010). Therefore, here we conclude that these phenomena may occur even in the absence of putative temporal correlations across time bins, pseudo-correlations caused by inappropriately assuming stationarity, or higher-order correlations, as previous studies have conjectured (Oizumi et al., 2009, 2010), due to the sole presence of destructive interference.

However, we found the emergence of destructive interference puzzling for at least two reasons. First, it never occurs when studying response aspects such as spike counts or latencies (Eyherabide, 2016), or when studying correlation importance using ΔI^D , as we have here shown (Section 3.7). Second, should ΔI^{DL} be exact as currently believed,

noise correlations would seemingly grow in importance with the number of populations, even if they transmit independent information.

These reasons notwithstanding, we cannot rigorously conclude that they imply flaws in ΔI^{DL} for at least two other reasons. First, whether or not ignoring response aspects differs from ignoring response probabilities remains unsettled (Eyherabide and Samengo, 2013; Latham and Roudi, 2013; Nirenberg and Latham, 2003; Schneidman et al., 2003). Second, although often overlooked in previous studies (Eyherabide and Samengo, 2013; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2009, 2010), ΔI^D is based on a different notion of information than ΔI^{DL} (Latham and Roudi, 2013; Nirenberg et al., 2001a; Nirenberg and Latham, 2003) and may overestimate the communication information loss (Eyherabide and Samengo, 2013; Latham and Nirenberg, 2005; Oizumi et al., 2009, 2010). Instead, we hypothesized that the phenomenon of destructive interference may only indicate that the intuition gained from traditional information theory applied to ignoring aspects of neural responses should be observed with caution when applied to ignoring aspects of response probabilities.

Unnoticed by previous studies is the fact that ΔI^{DL} is based on only one possible construction of optimal NI decoders, that we called joint NI decoding, which identifies all stimulus features simultaneously (Fig. 8(a)). However, here we showed that optimal NI decoders can be constructed in other ways that can potentially outperform the state of the art. In particular, and despite previous controversies, here we proved that optimal NI decoders constructed to identify independent stimulus features separately but in parallel can completely overcome the destructive interference (Fig. 8(b)).

This finding was puzzling because, when information is independent, joint and parallel NI decoders typically produce the same estimates. This observation seemingly rules out construction differences as the ultimate cause of destructive interference, but provided us with valuable insight into potential flaws of ΔI^{DL} , even when information is not independent. Specifically, we hypothesized that parallel NI decoders seemingly outperform joint NI decoders because they seemingly require that ΔI^{DL} be computed with codebooks that preserve the independence of the stimulus features.

Our tests on this hypothesis revealed for the first time in neural coding that, contrary to what occurs when studying response aspects such as spike counts or latencies, the average decoding error probability over the set of all codebooks used in the original derivation of ΔI^{DL} need not be representative of the average decoding error probability for smaller sets of codebooks when studying noise correlations. This result rigorously proved for the first time that, contrary to previously thought (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2016, 2009, 2010), ΔI^{DL} may overestimate the communication information loss caused by

ignoring noise correlations in optimal decoding even when stimuli cannot be perfectly identified or when information is not independent. Most importantly, this observation puts forward testing different sets of codebooks as one possible strategy for solving the overestimation.

We point out that the concept of destructive interference is not limited to the communication notion of information, can also be extended to the axiomatic notion of information. From our results on joint and parallel decoding, it follows immediately that this axiomatic destructive interference is not related to overestimations of the axiomatic information loss but to differences in tie-breaking rules adopted during the construction of joint and parallel NI decoders. Furthermore, it can be positive or negative, but the minimum axiomatic destructive interference over all possible conventions is always negative. Whether tie-breaking rules play any role in the communication destructive interference, or whether joint NI decoders can actually extract more communication information than parallel NI decoders, still remain open questions.

Since its introduction, the information-theoretical measure ΔI^{DL} has been deemed the exact information loss caused by ignoring noise correlations in optimal decoding (Ince et al., 2010; Latham and Nirenberg, 2005; Latham and Roudi, 2013; Oizumi et al., 2016, 2009, 2010). However, our results prove that ΔI^{DL} is biased and that the overestimation can reach $\approx 100\%$ of the encoded information. Hence, using ΔI^{DL} in its basic form may lead to wasting experimental and computational resources, which can be avoided by estimating the communication information loss as we propose here. These estimates close the gap between axiomatic and communication information losses, thereby opening up the possibility that traditional relations between information and decoding observed when studying response aspects such as spike counts and latencies are also valid when studying noise correlations. In practice, our results indicate that noise correlations need not be as necessary as previously thought, and may potentially contribute to reduce the cost and complexity of computational brain models and neuroprosthetics.

5 Conclusion

Assessing the role of noise correlations in neural decoding is fundamental, not only for understanding how the brain perform computations and turn them into perceptions, decisions and actions, but also for estimating the amount of resources and the level of complexity required to study brain function and to construct neural prosthetics. This study sheds new light into their role by revealing and resolving unforeseen limitations of an approach that, due to its rigorous information-theoretical foundations, has always been deemed exact. Our analysis was entirely conducted taking into account the fun-

damental differences between the notions of information associated with this and other approaches. In this way, we avoided the confounds of previous studies, and rigorously proved that the currently-deemed-exact approach overestimates the information loss caused by ignoring noise correlations in optimal decoding. In practice, our study shows that the cost of ignoring noise correlations for studying brain computations and information integration, when evaluated using the communication notion of information, can be much lower than currently thought, thereby potentially saving experimental and computational resources, and contributing to develop simpler and more efficient neuroprosthetics and technological applications.

Acknowledgments

This work was supported by the Academy of Finland, Centre of Excellence in Inverse Problems (project number 213476) and Computational Sciences Program (project number 135198).

References

- Abbott, L., Dayan, P., 1999. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 11 (1), 91–101.
- Aflalo, T., Kellis, S., Klaes, C., Lee, B., Shi, Y., Pejsa, K., Shanfield, K., Hayes-Jackson, S., Aisen, M., Heck, C., Liu, C., Andersen, R., 2015. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science* 348 (6237), 906–910.
- Akam, T., Kullmann, D., 2014. Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nat. Rev. Neurosci.* 15 (2), 111–122.
- Averbeck, B., Latham, P., Pouget, A., 2006. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7 (5), 358–366.
- Averbeck, B., Lee, D., 2006. Effects of noise correlations on information encoding and decoding. *J Neurophysiol* 95 (6), 3633–3644.
- Babiloni, F., Astolfi, L., 2014. Social neuroscience and hyperscanning techniques: past, present and future. *Neurosci Biobehav Rev* 44, 76–93.
- Bialek, W., 1987. Physical limits to sensation and perception. *Annu. Rev. Biophys. Biophys. Chem.* 16 (1), 455–478.

- Bialek, W., Rieke, F., de Ruyter van Steveninck, R., Warland, D., 1991. Reading a neural code. *Science* 252 (5014), 1854–1857.
- Bouton, C., Shaikhouni, A., Annetta, N., Bockbrader, M., Friedenberg, D., Nielson, D., Sharma, G., Sederberg, P., Glenn, B., Mysiw, W., Morgan, A., Deogaonkar, M., Rezai, A., 2016. Restoring cortical control of functional movement in a human with quadriplegia. *Nature*.
- Brenner, N., Strong, S., Koberle, R., Bialek, W., de Ruyter van Steveninck, R., 2000. Synergy in a neural code. *Neural Comput.* 12 (7), 1531–1552.
- Casella, G., Berger, R., 2002. *Statistical Inference*, second edition Edition. Duxbury.
- Cohen, M. R., Kohn, A., 2011. Measuring and interpreting neuronal correlations. *Nat Neurosci* 14 (7), 811–819.
- Cover, T., Thomas, J., 2006. *Elements of Information Theory*, second edition Edition. Wiley-interscience.
- Delis, I., Berret, B., Pozzo, T., Panzeri, S., 2013. A methodology for assessing the effect of correlations among muscle synergy activations on task-discriminating information. *Front Comput Neurosci* 7, 54.
- Duda, R., Hart, P., Stork, D., 2000. *Pattern Classification*, second edition Edition. Wiley John & Sons.
- Ernst, M., Banks, M., 2002. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415 (6870), 429–433.
- Eyherabide, H., Rokem, A., Herz, A., Samengo, I., 2008. Burst firing is a neural code in an insect auditory system. *Front. Comput. Neurosci.* 2 (3).
- Eyherabide, H., Samengo, I., 2010. Time and category information in pattern-based codes. *Front. Comput. Neurosci.* 4, 145.
- Eyherabide, H., Samengo, I., 2013. When and why noise correlations are important in neural decoding. *J. Neurosci.* 33 (45), 17921–17936.
- Eyherabide, H. G., 2016. Neural stochastic codes, encoding and decoding. Preprint. Available from: [arxiv:1611.05080](https://arxiv.org/abs/1611.05080). Cited 6 January 2017.
- Fano, R., 1961. *Transmission of Information*. The M.I.T. Press.
- Gallager, R., 1968. *Information theory and reliable communication*. John Wiley and Sons, Inc., New York, USA.

Gawne, T., Richmond, B., 1993. How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* 13 (7), 2758–2771.

Geisler, W., 2011. Contributions of ideal observer theory to vision research. *Vision Res.* 51 (7), 771–781.

Gross, J., Hoogenboom, N., Thut, G., Schyns, P., Panzeri, S., Belin, P., Garrod, S., 2013. Speech rhythms and multiplexed oscillatory sensory coding in the human brain. *PLoS Biol.* 11 (12), e1001752.

Hari, R., Kujala, M., 2009. Brain basis of human social interaction: from concepts to brain imaging. *Physiol Rev* 89, 453–479.

Harvey, M., Saal, H., Dammann III, J., Bensmaia, S., 2013. Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. *PLoS Biol.* 11 (5), e1001558.

Huk, A., 2012. Multiplexing in the primate motion pathway. *Vision Res.* 62, 173–180.

Ince, R., Senatore, R., Arabzadeh, E., Montani, F., Diamond, M., Panzeri, S., 2010. Information-theoretic methods for studying population codes. *Neural Netw.* 23 (6), 713–727.

Jaynes, E., 2003. *Probability Theory: the Logic of Science*. Cambridge university press.

Knill, D., Richards, W., 1996. *Perception as Bayesian Inference*. Cambridge University Press.

Landy, M., Maloney, L., Johnston, E., Young, M., 1995. Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Res.* 35 (3), 389–412.

Latham, P., Nirenberg, S., 2005. Synergy, redundancy, and independence in population codes, revisited. *J. Neurosci.* 25 (21), 5195–5206.

Latham, P., Roudi, Y., 2013. *Role of correlations in population coding*. CRC Press, Ch. 7, pp. 121–138.

Meister, M., Hosoya, T., 2001. Are retinal ganglion cells independent encoders? Preprint.

URL http://www.gatsby.ucl.ac.uk/~pel/critics/Meister_2001_Encoders.pdf

Meister, M., Lagnado, L., Baylor, D., 1995. Concerted signaling by retinal ganglion cells. *Science* 270 (5239), 1207–1210.

- Merhav, N., Kaplan, G., Lapidot, A., Shamai Shitz, S., 1994. On information rates for mismatched decoders 40 (6), 1953–1967.
- Meytlis, M., Nichols, Z., Nirenberg, S., 2012. Determining the role of correlated firing in large populations of neurons using white noise and natural scene stimuli. *Vision Res.* 70, 44–53.
- Montani, F., Kohn, A., Smith, M., Schultz, S., 2007. The role of correlations in direction and contrast coding in the primary visual cortex. *J Neurosci* 27 (9), 2338–2348.
- Nirenberg, S., Carcieri, S., Jacobs, A., Latham, P., 2001a. Retinal ganglion cells act largely as independent encoders. *Nature* 411 (6838), 698–701.
- Nirenberg, S., Carcieri, S., Jacobs, A., Latham, P., 2001b. Supplementary information. URL physiology.med.cornell.edu/faculty/nirenberg/lab/pel/critics/Meister
- Nirenberg, S., Latham, P., 1998. Population coding in the retina. *Curr. Opin. Neurobiol.* 8 (4), 488–493.
- Nirenberg, S., Latham, P., 2003. Decoding neuronal spike trains: how important are correlations? *Proc. Natl. Acad. Sci. U.S.A.* 100 (12), 7348–7353.
- Oizumi, M., Amari, S., Yanagawa, T., Fujii, N., Tsuchiya, N., 2016. Measuring integrated information from the decoding perspective. *PLoS Comput. Biol.* 12 (1), e1004654.
- Oizumi, M., Ishii, T., Ishibashi, K., Hosoya, T., Okada, M., 2009. A general framework for investigating how far the decoding process in the brain can be simplified. *Adv. Neural Inf. Process. Syst.*, 1225–1232.
- Oizumi, M., Ishii, T., Ishibashi, K., Hosoya, T., Okada, M., 2010. Mismatched decoding in the brain. *J. Neurosci.* 30 (13), 4815–4826.
- Oppenheim, A., Willsky, A., Nawab, S., 1997. *Signals and systems*, second edition Edition. Prentice Hall.
- Oram, M., Földiák, P., Perrett, D., Sengpiel, F., 1998. The 'ideal homunculus': decoding neural population signals. *Trends Neurosci.* 21 (6), 259–265.
- Panzeri, S., Brunel, N., Logothetis, N., Kayser, C., 2010. Sensory neural codes using multiplexed temporal scales. *Trends Neurosci.* 33 (3), 111–120.
- Panzeri, S., Golledge, H., Zheng, F., Tovée, M., Young, M., 2001. Objective assessment of the functional role of spike train correlations using information measures. *Vis. Cogn.* 8 (3-5), 531–547.

- Pereda, E., Quiroga, R., Bhattacharya, J., 2005. Nonlinear multivariate analysis of neurophysiological signals. *Prog. Neurobiol.* 77 (1–2), 1–37.
- Quiroga, R., Panzeri, S., 2009. Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.* 10 (3), 173–185.
- Rolls, E., Treves, A., 2011. The neuronal encoding of information in the brain. *Prog. Neurobiol.* 95 (3), 448–490.
- Samengo, I., 2002. Estimating probabilities from experimental frequencies. *Phys Rev E Stat Nonlin Soft Matter Phys* 65 (4), 046124.
- Samengo, I., Treves, A., 2000. Representational capacity of a set of independent neurons. *Phys. Rev. E* 63 (1), 011910.
- Schneidman, E., Bialek, W., Berry, M., 2003. Synergy, redundancy, and independence in population codes. *J. Neurosci.* 23 (37), 11539–11553.
- Schneidman, E., Puchalla, J., Segev, R., Harris, R., Bialek, W., Berry, M., 2011. Synergy from silence in a combinatorial neural code. *J. Neurosci.* 31 (44), 15732–15741.
- Shannon, C., Weaver, W., 1949. *The Mathematical Theory of Communication*. University of Illinois Press.
- Simoncelli, E., 2009. *Optimal estimation in sensory systems*. The MIT Press, pp. 525–535.
- Thomson, E., Kristan, W., 2005. Quantifying stimulus discriminability: a comparison of information theory and ideal observer analysis. *Neural Computation* 17 (4), 741–778.
- Warland, D., Reinagel, P., Meister, M., 1997. Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* 78 (5), 2336–2350.
- Womelsdorf, T., Lima, B., Vinck, M., Oostenveld, R., Singer, W., Neuenschwander, S., Fries, P., 2012. Orientation selectivity and noise correlation in awake monkey area v1 are modulated by the gamma cycle. *Proc Natl Acad Sci U S A* 109 (11), 4302–4307.
- Woodward, P., Davies, I., 1952. Information theory and inverse probability in telecommunication. *Proc. I. E. E.* 99 (58), 37–44.
- Zhang, W., Chen, A., Rasch, M., Wu, S., 2016. Decentralized multisensory information integration in neural systems. *J. Neurosci.* 36 (2), 532–547.